NATIONAL
URANIUM
TAILINGS
PROGRAM

[illegible]

Canadä

APPROACHES TO RISK ASSESSMENT FOR CANADIAN URANIUM MILL TAILINGS
PHASE II

DSS File #14SQ.23241-5-1717
Contract No. 0SQ85-00225

A Research Report prepared for:

The National Uranium Tailings Program
Canada Centre for Mineral and Energy Technology
Energy, Mines and Resources Canada
555 Booth Street
Ottawa, Ontario
K1A 0G1

by

SENES Consultants Limited
499 McNicoll Avenue
Willowdale, Ontario
M2H 2C9

March 1986

SEP 8 1986
MICROMEDIA

176 pp

ABSTRACT

This study represents the second phase of work directed at developing a methodology for quantifying the risks associated with uranium mill tailings. We adopt the concept of risk as a "set of triplets"; consisting of scenarios, the likelihoods of scenarios occurring and their consequences, should they occur. We recognize that both the likelihood of an event occurring and its potential consequences may be very uncertain quantities. The methodology developed and applied in this study allows one to incorporate these uncertainties in quantifying risk. The results are presented in the form of "level 2" risk curves.

The study also investigates a number of other matters related to the quantification of risk. First, several possible applications of the de minimis concept in probabilistic risk analysis are identified and discussed. Secondly, the aspect of considering time in risk analyses is examined. An approach to incorporating time varying consequences is developed and demonstrated. Finally, we briefly discuss the presentation and interpretation of results and the appropriate role of quantitative risk analysis in risk management decision-making.

ACKNOWLEDGEMENTS

Many people contributed to this study. The study was carried out under the direction of Mr. R. A. Knapp of SENES. Principal technical contributors from SENES include Dr. D. B. Chambers, Mr. M. L. Murray and Mr. D. T. Ellenor.

Dr. Stanley Kaplan developed the approach for incorporating uncertainty into the probabilistic risk assessment and was the source of much useful insight into the risk assessment process. Dr. Kaplan's effort in this study is very much appreciated.

Dr. Granger Morgan was kind enough to take time from his hectic schedule to provide a much needed perspective on risk assessment in decision making.

Dr. Park Reilly provided sage counsel concerning various subtleties of statistical inference and probability. His advice contributed substantially to the successfull completion of this project and is gratefully acknowledged.

Finally, we wish to thank Mr. Ron Holmes of the NUTP for his patience and many useful suggestions throughout the course of the study.

EXECUTIVE SUMMARY

This study, carried out as part of the National Uranium Tailings Program (NUTP), was conducted to extend and refine the risk analysis methodology developed during a previous study. The methodology adopts the point of view that the idea of risk can be defined quantitatively as a set of triplets; expressing the possible scenarios, their likelihoods, and their consequences. In the present application, a "scenario: is taken to be an unfolding of the entire time history of a tailings pile over a 1,000 year period. The "consequences" of these scenarios are the long-term impacts on humans and the environment arising from the tailings. The evaluation of the consequences requires the use of mathematical models to simulate the movement of these substances in the tailings pile and in the environment.

The previous study assumed that event likelihoods were known with certainty and used mean value estimates for scenario consequences. That methodology was modified to explicitly incorporate the uncertainty associated both with event likelihoods and scenario consequences. This essentially involved generating a large number of "level 1" risk curves (which ignore uncertainty) and summarizing them via "level 2" risk curves, which illustrate the uncertainty in the risk analysis.

A further objective of the study was to illustrate the inclusion of time-dependency in a risk analysis. The effects of the tailings (e.g. radiation dose to a nearby receptor) change with time even in the absence of such episodic events as floods or forest fires. This time dependent behaviour was incorporated by dividing the 1000 year simulation period into several intervals, over which the consequences did not vary greatly. The likelihoods and consequences of events occurring in each time window were then incorporated into the analysis.

The study also examined other matters related to the quantification of risk. First, possible applications of the de minimis concept in probabilistic risk analysis were identified and discussed. The presentation and interpretation of results and the appropriate role of quantitative risk analysis in risk management decision-making was also examined.

The study concluded that the methodology represents a viable technique for incorporating time-dependence and uncertainty into risk analysis. Although it provides a powerful tool for combining the consequences of several events (in proportion to their likelihood) considerable effort is required to perform the analysis properly. A risk analysis, such as that outlined in this report, may be most useful in providing insight into a problem, rather than in generating "the answer".

TABLE OF CONTENTS

Page

Abstract i
Acknowledgement ii
Executive Summary S-1

1.0 INTRODUCTION

1.1 Background 1-1
1.2 Study Objectives 1-2
1.3 Review of Relevant Concepts 1-2

2.0 SELECTED CONSIDERATIONS IN PROBABILISTIC RISK ASSESSMENT

2.1 Overview 2-1
2.2 Application of the *De Minimis* Concept 2-1
2.2.1 The *De Minimis* Concept 2-1
2.2.2 Role of the *De Minimis* Concept in Uranium Tailings Management 2-3
2.2.3 Application to Probabilistic Risk Assessment 2-5
2.3 Incorporation of Time Dependence 2-8
2.3.1 Source Term Characteristics 2-9
2.3.2 Frequency of Natural Events 2-9
2.3.3 Timing of Event Occurrence 2-10
2.3.4 Competency of Engineered Structures 2-11
2.3.5 Population Shifts 2-11
2.4 Incorporation of Uncertainty 2-12

3.0 SITE DESCRIPTION AND DEFINITION OF EVENTS, SCENARIOS AND CONSEQUENCES

3.1 Description of the Reference Facility 3-1
3.2 Events Considered 3-2
3.3 Reference Receptor 3-4
3.4 Complicating Factors Considered 3-4

Table of Contents (continued)

Page

4.0 DEMONSTRATION OF A METHODOLOGY FOR PROBABILISTIC RISK ASSESSMENT

4.1 A Review of Relevant Concepts - Continued 4-1

4.2 Example 1: Incorporating Uncertainty 4-3

4.2.1 Uncertainty in Scenario Frequencies 4-4

4.2.2 Uncertainty in Scenario Consequence 4-5

4.2.3 Combining Uncertainties 4-6

4.2.4 Summary of Methodology 4-9

4.3 Example 2: Incorporating Time Variance 4-10

4.4 Example 3: A Non-Radiological Consequence 4-14

4.5 Example 4: A Cumulative Consequence 4-15

5.0 PRESENTATION AND INTERPRETATION OF RESULTS 5-1

5.1 Communicating Probabilistic Results 5-2

5.2 Interpreting Risk Results 5-5

5.3 The Role of Quantitative Risk Analysis in Risk Management Decision-Making 5-8

6.0 SUMMARY AND CONCLUSIONS

6.1 Application of the *De Minimis* Concept 6-1

6.2 Incorporation of Time Dependence 6-2

6.3 Incorporation of Uncertainty 6-3

6.4 Presentation and Interpretation of Results 6-3

6.5 Conclusions/Recommendations 6-4

7.0 REFERENCES 7-1

1.0 INTRODUCTION

1.1 Background

In 1982 Energy, Mines and Resources Canada established the National Uranium Tailings Program (NUTP) which is intended to pull together previous work on Canadian uranium mill tailings, to fund additional work where significant information gaps are found and to develop the capability to predict the long-term effects of tailings. Such evaluations require the use of mathematical models to simulate the behaviour of the real system(s).

Environmental models are, however, only partial representations of real systems and contain a number of parameters which are either only imperfectly known or are inherently variable. Consequently, the results of all such calculations are uncertain. Uncertainties also arise when these models are used to predict effects over conditions and periods of time different from those for which the models and model parameters were developed. Probabilistic analysis allows us to quantify some of the uncertainty associated with the long-term effects of uranium mill tailings and to generate distributions of expected consequences (e.g. the distribution of maximum annual dose to the most exposed receptor). Typically, such an analysis involves the simulation of one model site, or event, at a time.

In 1985, a risk study (hereafter referred to as Phase I) was funded by the NUTP to help make the predictions from such models more useful in decision making. The study identified and applied a methodology for combining the predicted consequences of many possible events to produce an overall estimate of the risk associated with a given facility or site. This methodology adopted the idea that risk can be considered as a set of triplets, each triplet being composed of a scenario (a particular combination of possible events), the likelihood of the scenario occurring and the consequence associated with the scenario. The results of these analyses were displayed using event trees or as "risk curves" as outlined in a subsequent section.

1.2 Study Objectives

This study had a number of objectives. The primary focus was on the development and demonstration of a methodology for incorporating time variation and uncertainty into estimation of the risk associated with uranium mill tailings. Several other issues were also examined. The general study objectives are summarized below.

As a follow-up to the work completed in the Phase I risk study, the following areas were to be examined:

1. Application of the *de minimis* concept to probabilistic risk assessments.
2. Further incorporation of time dependency into risk analysis.
3. Incorporation of uncertainties about the annual frequency of occurrence of events and about the predicted consequences.
4. Interpretation and presentation of risk assessment results in helping to make decisions.
5. Application of the concepts to a uranium tailings facility.

1.3 Review of Relevant Concepts

Risk as a Set of Triplets

As in a previous study (SENES, 1985a), this study adopted a working definition of risk as a set of triplets (Kaplan and Garrick, 1981). Risk can be represented as the answers to the three questions:

1. What can happen? (i.e. what can go wrong?)
2. How likely is that to happen?
3. If it does happen what are the consequences? (i.e. what are the damages?)

The answers may be conveniently written in the form of a triplet:

FIGURE 1·1

RISK AS A EVENT TREE

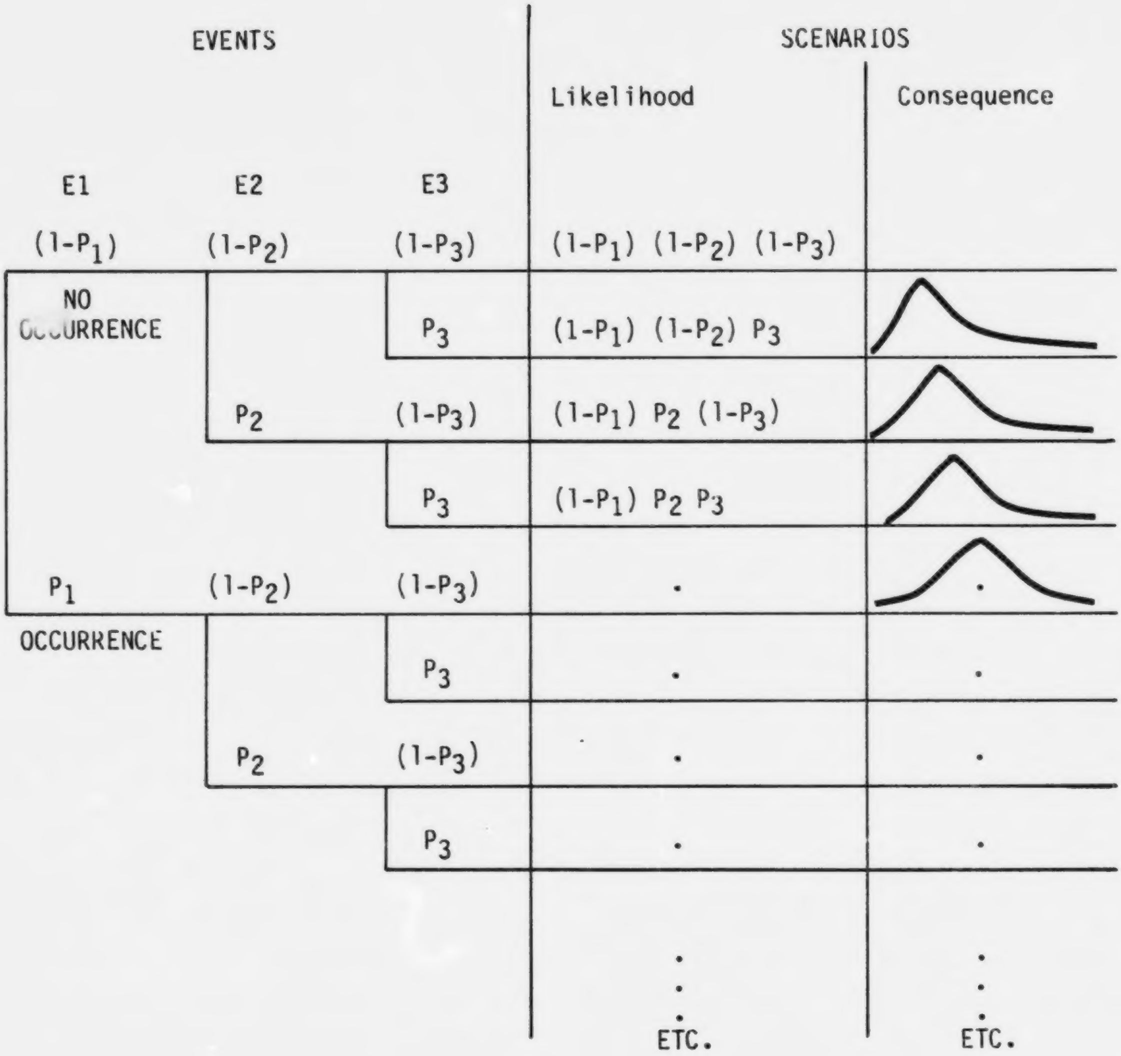
EVENTS
SCENARIOS
Likelihood
Consequence
E1
E2
E3
$(1-P_1)$
$(1-P_2)$
$(1-P_3)$
$(1-P_1)\ (1-P_2)\ (1-P_3)$
NO OCCURRENCE
P_3
$(1-P_1)\ (1-P_2)\ P_3$
P_2
$(1-P_3)$
$(1-P_1)\ P_2\ (1-P_3)$
P_3
$(1-P_1)\ P_2\ P_3$
P_1
$(1-P_2)$
$(1-P_3)$
OCCURRENCE
P_3
P_2
$(1-P_3)$
P_3
ETC.
ETC.

Note:

P_1, P_2, P_3 are the probabilities of events E_1, E_2, E_3 occuring one or more times over the simulation period.

$$<S_i, l_i, x_i>$$

where S_i is a name or a description of a "scenario" or what can happen; l_i is the likelihood of that scenario occurring; and x_i is a "damage index" or some measure of the consequences of the scenario.

More formally, using braces $\{ \}$ to denote "set of", the risk R is defined as:

$$R = \{<S_i, l_i, x_i>\} \quad (1.1)$$

Thus, for analytic purposes, the risk "is" the set of triplets in equation 1.1.

Event Trees

The information contained in the set of triplets may be summarized in the form of an event tree, an example of which is presented in Figure 1.1. In this case there are only three basic events that can occur: E_1, E_2, and E_3. From these events,it is possible to construct "scenarios", which are simply various combinations of the basic events which can occur during the period of interest - the simulation period. If we "know" the annual frequency of occurrence of E_1, E_2 and E_3 (i.e. the chance of E_1, E_2, E_3 occurring in any year) we can determine the probability of non-occurrence $(1-P_i)$ and the probability of one or more occurrences (P_i) over the simulation period. From these we can determine the likelihood of occurrence for each scenario, as indicated on Figure 1.1.

Two further points are worth noting. First, it is common to structure the event tree such that the largest consequences are at the bottom of the event tree and the smallest at the top. In most situations the ordering is apparent from the nature of the consequence being examined. The worked examples provided later in this report should help to clarify this point. Second, each scenario actually represents a class of scenarios and could, if desired (and with additional information), be further subdivided into finer categories. Alternatively, scenarios which differ very little in consequence can be combined to form a smaller event tree. Overall, an event tree is simply a convenient way to represent the various scenarios, likelihoods and consequences that constitute the "risk".

Mathematical models are used to assess the consequences of the scenarios. However, as previously noted, models are only partial and imperfect reflections of real systems and contain a number of parameters which are either only imperfectly known or inherently variable. Probabilistic analysis allows us to incorporate the uncertainty associated with these input parameters. In a probabilistic analysis, the input parameters are specified as distributions, or ranges of values, rather than as single numbers. The models are then evaluated many times with different values selected for each (uncertain) input parameter on each trial. The result of such an analysis is a distribution of results, rather than a single value estimate. The output from a probabilistic analysis. (i.e. the distribution of calculated values) may be interpreted as a subjective probability describing in quantitative terms what we are entitled to believe about the value of the output variable (i.e. the consequence). The curves shown on Figure 1.1 under the title "consequence" represent output distributions for each scenario.

Probabilistic techniques have been applied in several previous studies of uranium mill tailings (SENES, 1984; SENES, 1985a,b). This report assumes a working knowledge of probabilistic analysis concepts. For background, or further information the reader is referred to any of the previous studies.

Level 1 Risk Curve

An alternative way to express "risk" is in the form of a risk curve such as that shown in Figure 1.2. Risk curves were constructed in the Phase I risk study in the following way. First, a particular attribute of the distribution of consequence, namely the mean conseqence $\bar{C}$, was used to represent the consequence of a particular scenario. Likelihoods were calculated for the

FIGURE 1·2
"RISK" AS A
CONSEQUENCE PROBABILITY CURVE

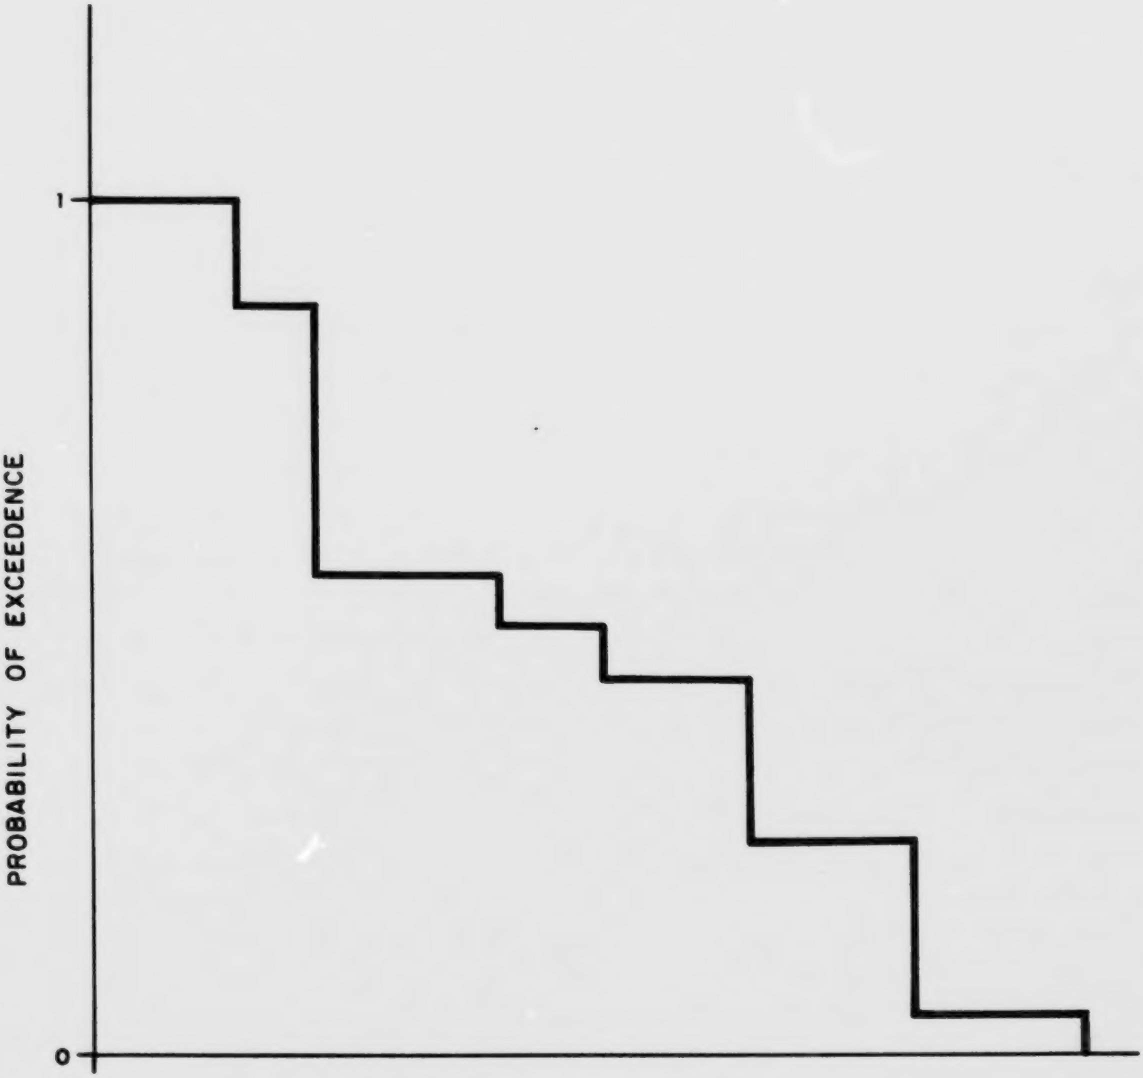

various scenarios, assuming that the annual frequency of occurrence of the individual events were precisely known. The risk curve was then constructed by summing likelihoods from bottom to top in the event tree (i.e. high consequence to low consequence). This produces a downward or complementary cumulative probability function such as that shown in Figure 1.2. The complementary cumulative probability is shown on the vertical axis and the consequence on the horizontal axis. The ordinate represents the cumulative probability that during the simulation period the mean consequence will be equal to or greater than the corresponding consequence indicated on the abscissa. Kaplan and Garrick (1981) refer to such curves as Level 1 risk curves. Later in this report the concept of communicating risk by a family of risk curves will be introduced.

Statistical Terminology

From the foregoing brief discussion it is already clear that there are some statistical words, namely frequency, probability and likelihood which are used in a special sense in this report and therefore must be carefully defined. For example, "likelihood" is used in a non-technical sense. As a technical word in statistics it has a very special meaning, but here it is to be taken as some yet-undefined measure of the chance of a particular occurrence (i.e. how likely is a particular scenario?)

The usages of "frequency" and "probability" are best described by a simple example. Suppose we are interested in the fraction of thumb tacks of a certain type which will land point up when tossed on a table. This is easily recognized as a binomial (or Bernoullian) situation in which the parameter p describes the fraction of tacks which will land point up when an indefinitely large number of identical tacks is tossed, or if a single tack is tossed an indefinitely large number of times.

Suppose that we know p and that we toss a set of n tacks an indefinitely large number of times. The fraction of this large number of tosses which will yield exactly x point up of the n tacks tossed in the set is given by the binomial formula.

$$\binom{n}{x} p^{x} \quad (1-p)^{n-x} \qquad (1)$$

Here we have twice used the concept of the observed fraction of outcomes of a certain type in an indefinitely large number of trials: once in defining p, and once in interpreting the bionomial formula. In this report both of these usages are called frequencies. We use the word to mean what most statiticians would call objective probability, long term relative frequency, or most likely just probability. Our use of the word frequency emphasizes the fact that it refers to the oucome of some physical event and we mean to imply that, at least in principle, these frequencies can always be measured simply by performing many, many trials and counting outcomes of the appropriate types.

Suppose now that we do not know p but are interested in estimating it. In an actual experiment a box of 25 tacks was shaken up and emptied on a table with the result that 14 of them landed point up. Then by standard Bayesian methods (see Box and Tiao 1973) we may deduce a posterior probability density function, in this case a Beta distribution function, which tells us what we may infer about p. It is plotted in Figure 1.3 and has the mathematical form.

$$Df(p) = \frac{\Gamma(26)}{\Gamma(14.5)\,\Gamma(11.5)} p^{13.5} (1-p)^{11.5} \qquad (2)$$

This describes a distribution of what we call probability in this report. Areas under the bell-shaped curve are proportional to the strength of belief which we are entitled to hold about the corresponding values of p. Thus we may say in the context of this report that the curve shown in Figure 1.3 describes probabilities of the various possible values of a frequency.

The essential point is that, in our usage, frequency describes a proportion of outcomes of an experiment which is, at least in principle, physically realizable; probability tells us how strongly we feel or are entitled to feel

FIGURE 1·3

POSTERIOR PROBABILITY DENSITY FUNCTION
FOR FRACTION OF TACKS (p)
WHICH LAND POINT UP

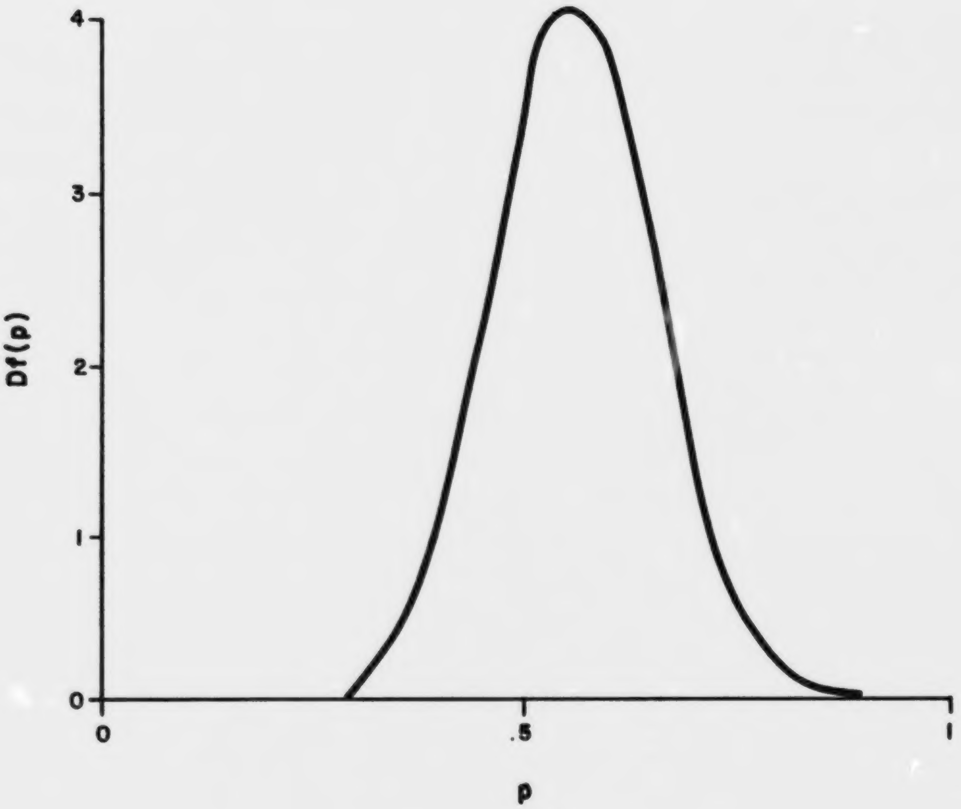

NOTE:

$$Df(p) = \frac{\Gamma(26)}{\Gamma(14.5)\Gamma(11.5)}\, p^{13.5}(1-p)^{11.5}$$

about different values of a parameter. The parameter itself may or may not be a frequency.

In standard books on statistics equation (2) would likely be described as a distribution of subjective probability concerning the possible values of the parameter p, which happens in this case to be an objective probability. However, many statistical writers do not emphasize (or possibly even recognize) the distinction between objective and subjective probability and would simply describe as probability both what we call probability and what we call frequency.

In actual practice the situation is usually not as simple as in the example. Most probabilisitc quantities fall somewhere between the extremes which we have called probability and frequency. For example suppose we have past data on wind speed and direction for a certain location. From these we could certainly produce a distibution of what we call frequency to describe them. However, as soon as we attempt to use these distributions to describe future wind speeds or directions, we introduce judgement as to how applicable past data, however good, are to the future. Thus there is necessarily some component of what we call probability in this usage. Recognizing this, we still retain both terms to indicate, at least roughly, the nature of the concepts which we describe.

2.0 SELECTED CONSIDERATIONS IN PROBABILISTIC RISK ASSESSMENTS

2.1 Overview

A previous study (SENES 1985a) describes a risk assessment methodology which can be used to evaluate risks arising from uranium mill tailings. The basic concepts of this approach have been outlined in Chapter 1. This chapter discusses, in a qualitative fashion, several extensions to the basic methodology.

One such issue involves the de minimis concept and its potential application to uranium mill tailings. In this chapter, the possible role of the de minimis concept in a probabilistic risk analysis of uranium tailings is discussed. In addition, two further topics are introduced, namely: the implications of incorporating time-dependency; and potential complications associated with including an estimate of uncertainty in both the likelihoods and consequences of the various scenarios. Examples illustrating the essential features of including time dependency and uncertainty in the risk methodology are provided in Chapter 4.

2.2 Application of the De Minimis Concept

2.2.1 The De Minimis Concept

The term *de minimis* has its origins in several Latin phrases that are often used to explain points of law. One example is the phrase *de minimis non curat lex* which can be translated as "the law does not concern itself with trifles." The essence of the *de minimis* concept is that there are conditions that are perceived to be insignificant because the risks they present are so small. In uranium tailings management, the *de minimis* concept might be a valuable guide in evaluating the acceptability of management techniques and assessing the potential effects imposed on man and the environment.

Since the supply of very small risks is virtually inexhaustible, and since resources available to manage such risks are always finite, individuals and organizations are constantly applying the principle. In very many instances, the de minimis concept is applied in an implicit rather than an explicit manner. That society has decided to tolerate exposures to various substances (or recognized the unavoidability of some exposure) is evident in several types of legislation such as drinking water criteria, air quality criteria, and food and drug regulations. Radiation protection agencies in Canada and the United States have used the essence of the de minimis concept by ruling that certain low level radioactive sources such as scintillation fluids and ionization-chamber smoke detectors can be placed in unlicensed waste disposal facilities. In 1985 the Atomic Energy Control Board issued a consultative document that set out their position on how the de minimis concept could be used as the basis for exempting certain radioactive materials, containing below de minimis levels of radioactivity from the standard licensing process (AECB, 1985). Even more recently, the United States Nuclear Regulatory Commission (USNRC) has recognized the desirability of a de minimis feature in radiation protection standards to avoid extending regulatory action beyond what is needed to adequately protect public health (USNRC, 1986).

The prime advantage often attributed to the de minimis concept is that it can help to prevent the spending of society's resources (which are finite) on insignificant matters to the detriment of those more deserving. The concept can also be used to circumvent the assessment of effects so small or improbable that the uncertainty inherent in the assessment techniques being used makes the estimates of the results unreliable. Similarly, the application of the concept has been cited as a way to rationalize the regulatory process by making it more consistent, cost effective, equitable and predictable.

The de minimis concept appears to offer several opportunities for improving the effectiveness of health protection or risk regulation. Three examples of potential economic benefits include: the elimination of wasteful practices in handling, shipping and disposing of low-level waste material; more effective use of special disposal repositories; and relief to regulators and industry from the monitoring, record keeping and assessment of trivial matters. The net results could be savings to the generators of de minimis materials and their shareholders, potential savings to their customers, the ability of regulators to shift resources to more pressing matters, and savings to the source of regulatory funding, the public.

Many of the considerations relevant to the philosophy, administration, and economic issues of health protection are also relevant to the ethical issues. A strong case can be made for stating that it can never be morally obligatory to waste resources, and that below certain levels of risk, the benefits to be derived from further attention to certain conditions or concentrations are so minute that expending resources on those conditions or concentrations would be wasteful, regardless of the way benefits are defined. Often the costs of even investigating a substance would be greater than the maximum potential benefits derived from doing something about it. This and many other considerations support the conclusion that adopting a de minimis concept is morally justified if correctly defined.

2.2.2 Role of the De Minimis Concept in Uranium Tailings Management

The de minimis concept is often presented as a component in the overall approach to regulating exposures or as a way of truncating or limiting dose in collective dose calculations as illustrated in Figure 2.1. The concept is that once an assessment has demonstrated that dose from a practice meets the maximum permissible criterion, then the need for further dose reduction could be based on the application of the ALARA (as low as reasonably achievable social and economic factors taken into account) principle for dose reductions between the maximum limit and a selected de minimis value.

The de minimis level represents a practical end point or "floor" for the range of application of ALARA analysis. It is important to note that a de minimis level should not be perceived to be the ultimate goal of exposure reduction (that is, the achievement of de minimis levels is not the objective of an ALARA analysis) but rather should be used only to identify exposures that need

FIGURE 2·1

RELATIONSHIP BETWEEN COMPONENTS OF RADIATION EXPOSURE REGULATION AND EXPOSURE LEVEL

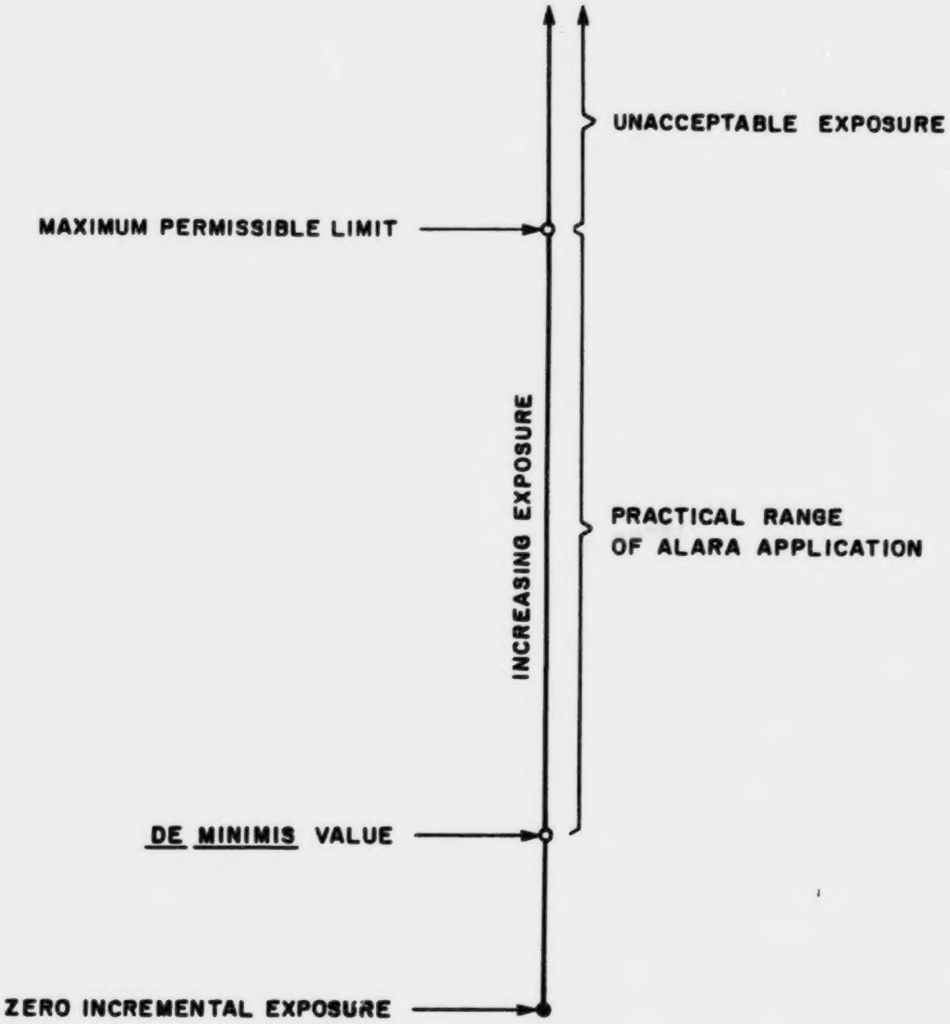
UNACCEPTABLE EXPOSURE
MAXIMUM PERMISSIBLE LIMIT
INCREASING EXPOSURE
PRACTICAL RANGE OF ALARA APPLICATION
DE MINIMIS VALUE
ZERO INCREMENTAL EXPOSURE

not be considered further.

A second role for the de minimis concept is to identify limits for the sizes of populations or durations of time over which collective doses would need to be considered. Supporters of this role argue that this would eliminate the integrations over space and time that have been used to predict large numbers of health effects (on a statistical basis), but which actually represent upper limits and reveal very little about real situations other than that the actual effects lie somewhere between zero and the numbers calculated. Conversely, this use of the concept has been criticized as not taking into account potentially important components of the collective dose and possibly ignoring several trivial exposures which could combine to produce non-trivial effects.

The recent standards for radiation protection proposed by the USNRC (USNRC, 1986) set out an approach to de minimis and collective dose evaluation similar to the approach suggested in the foregoing. According to the USNRC the de minimis level would be a lower-limit for regulatory concern. The USNRC are proposing de minimis values for both the most exposed member of the public and as a cut-off for collective dose calculations.

Based on these considerations, two general applications are identified for the de minimis concept in the management of uranium tailings. The first role is to evaluate specific management practices by comparing the dose to the critical receptor to the de minimis value. If the receptor's dose is less than the de minimis value, no further health protection effort is needed. The second application is for the de minimis value to identify the point of truncation, either in time or space, for population or collective dose or risk calculations. Both of these applications help to place limits on the efforts used in ALARA-type calculations which might be used to compare alternative tailings close-out options.

As discussed elsewhere (SENES, 1985c), one possible form that the application of the de minimis concept to uranium tailings management could take is as follows:

- de minimis levels of risk would be defined for individual dose and for truncating collective dose estimates, the former being the greater of the two since the possibility of exposure to more than one source would be greater for a group than an individual;
- both de minimis risk levels would be translated to exposures or doses;
- to evaluate a specific practice, all pertinent pathways would be analysed for the critical receptor and if the exposures were less than the de minimis value for the individual, no further effort would be needed for exposure reduction;
- if pathways analysis indicated that exposures greater than the individual de minimis level could be anticipated, the second de minimis value would be used to identify the appropriate point of truncation in time or space for estimates of collective dose.

This process could be applied in principle to either radioactive or non-radioactive releases.

2.2.3 Application to Probabilistic Risk Assessment

The practical use of the de minimis concept requires re-evaluation in terms of a probabilistic risk assessment. One of the objectives of probabilistic modelling is to provide a more complete picture of a consequence than is possible with the conventional deterministic analysis. The output of a probabilistic analysis is generally presented as a probability distribution of consequence. Because the result is a distribution and not a single value, the possible application of the de minimis concept as discussed previously as a limiting dose for an individual or for truncating collective dose is no longer as simple a task.

If the de minimis concept were applied as the lower bound of regulatory concern then it could be argued that any probabilistic trials that produce dose values less than the de minimis level would not be included in the risk assessment. The de minimis value would be interpreted as the "0" value on the abscissa when plotting the risk curve. Trials with doses below the de minimis level simply would be excluded. This concept is illustrated in Figure 2.2. By excluding values below de minimis we are, in effect, saying such low values are of no interest, and need not be considered in any circumstance. In effect, in calculating attributes of dose (such as mean dose) this assumption is equivalent to assigning all dose calculated to be smaller than the de minimis dose, a value of "zero". The "mean" dose calculated in this way can never exceed the mean dose calculated in the regular way viz:

$$\frac{\Sigma D_i (D_i < \underline{de}\ \underline{minimis} = 0)}{N} \leq \frac{\Sigma D_i}{N}$$

In calculating both mean doses (including and excluding de minimis values), one divides the sum of the valid observations (D_i) by the number of trials (N). That is, by setting certain observations equal to zero, we do not reduce the number of trials (N) used in calculating the mean.

In any situation where all calculated values were smaller than the de minimis value, the mean dose would itself be considered de minimis or trivial or insignificant. Whether or not such an approach is appropriate is a matter of philosophical interest.

Alternatively, the de minimis level could be used as a "marker" for comparing options. For example, if de minimis doses can truly be considered as trivial then one would consider giving more weight to options that had a higher percentage (i.e. a greater probability) of the calculated dose values below the de minimis value. This concept is not totally appropriate because the shape of the distribution is a function of both the system and of our state of knowledge about the system. In some situations, the spread of values is narrow while for other systems the spread may be much wider. Consider Figure 2.3. If option 1 was well defined, the mean dose could be well above the de minimis level with little of the curve being lower than the de minimis level. On the other hand, option 2, which is poorly defined, has the same mean dose but a larger portion of the curve is below the de minimis level. If this were the only criterion, we might conclude this is a better option. However, option 2 could also have much higher doses (although of low probability) which thus

FIGURE 2-2

IMPLICATION OF DE MINIMIS
EXCLUDING DE MINIMIS VALUES

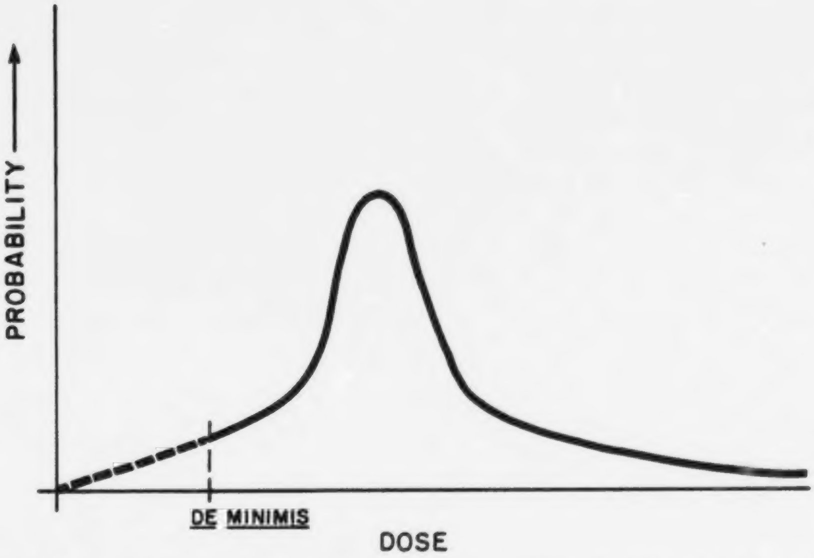

A) PROBABILITY DISTRIBUTION WITH EXCLUSION OF DE MINIMIS VALUES

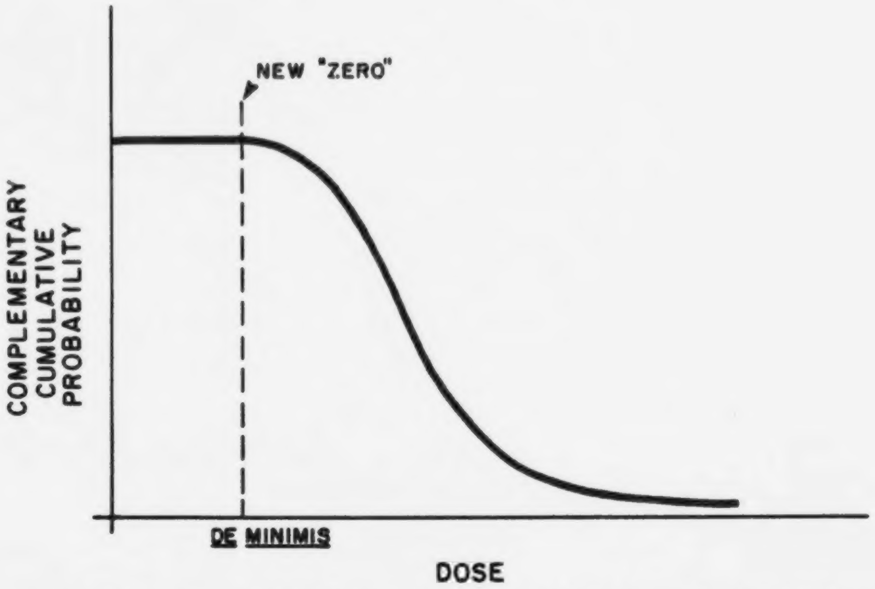

B) RISK CURVE SUGGESTS THERE IS NO RISK BELOW DE MINIMIS LEVEL

FIGURE 2·3

IMPLICATION OF *DE MINIMIS* FOR COMPARING OPTIONS

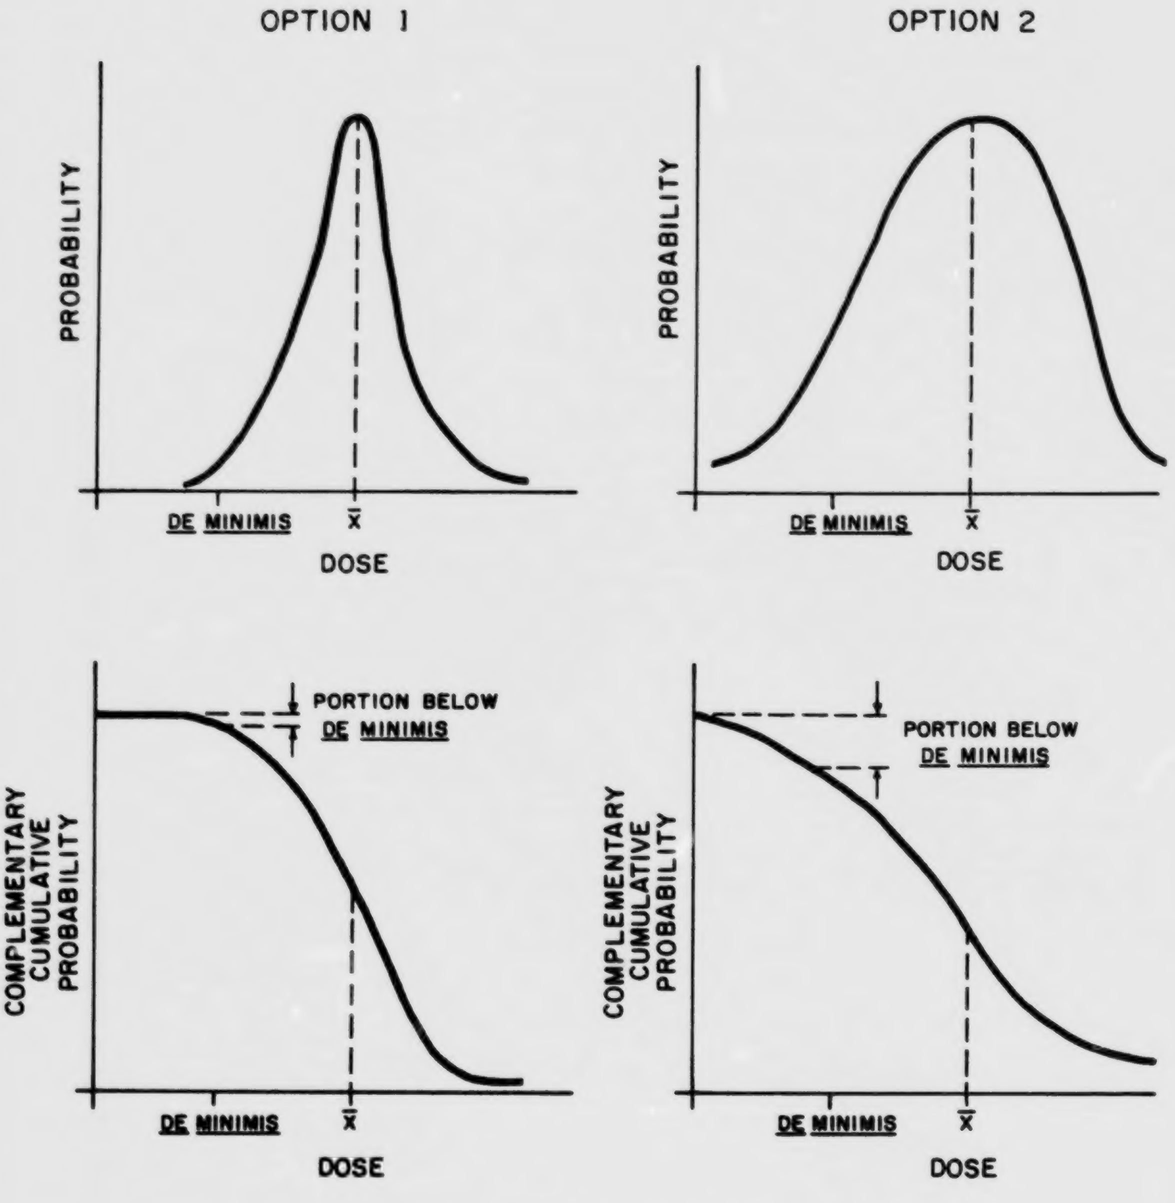

NOTE: BOTH OPTIONS HAVE THE SAME MEAN DOSE $\bar{X}$

could offset the potential benefits of having a large component of the dose below the de minimis level. If the mean dose were estimated to be below the de minimis level, the preferred option would seem to be the option for which we are most confident that the mean dose is actually below the de minimis level. This example illustrates that one component of the probabilistic curve cannot be viewed in isolation from the rest of the curve or indeed from other factors which might, with appropriate weighting factors, be considered in a multi-criteria approach.

The possible use of the de minimis concept for truncating dose in time and space is also complex. For a deterministic assessment, one would simply find the point in time or space where the calculated dose was less than the de minimis level and truncate collective dose calculations at that point. A similar approach is feasible with a probabilistic assessment. On each probabilistic trial one could continue to calculate individual dose in time and space until the calculated value was less than the de minimis value. All calculated values up to this point would be used in the calculation of collective dose for that particular trial. Thus, truncation would occur at different points in time and space for different trials. The drawback to this approach is that the times and/or distances required to reach de minimis levels might be very large for a few trials. In these instances, collective dose calculations could be truncated for other reasons, such as limitations of the simulation method.

Alternatively, one might calculate distributions of individual dose at many points in time and space. Different attributes (eg. mean, median) of these distributions would then be compared with the de minimis value. Those points in time and space having distributions of dose with (for example) mean values below de minimis would be excluded from the calculation of collective dose. One might also wish to specify a certain statistical confidence that the mean value is actually below de minimis. The problem with this approach is that individual doses well above the de minimis value might exist even for those distributions whose means fall below de minimis.

Our view is that de minimis remains a potentially useful concept for truncation of collective dose in time and space. There seem to be two basic approaches for application of the concept. The first approach might reject the portion of the output distributions which fall below the de minimis value. The second approach might be to reject or accept an entire output distribution based on the relation between some attribute of the distribution (eg. the mean) and the de minimis value. For example, the truncation dose could be set at either 1/10 the de minimis level or at the point where there is specified confidence, say 95%, that the mean calculated dose to a particular receptor will not exceed the de minimis value. The first approach of selecting "1/10" de minimis has been suggested by the AECB (Chatterjee, et al., 1982) as a method to be used in ALARA-type calculations to account for multiple sources and the variations in people's response to radiation. The second is an alternative that is both simple, and perhaps better utilizes the probabilistic modelling results. Assuming that the population can be defined it would be appropriate to include all of the dose calculated, including that portion below de minimis, for cumulative dose evaluations. These concepts are illustrated in Figure 2.4.

2.3 Incorporation of Time Dependence

There are many time-variant factors which can have significant implications in risk analysis. Some of the more important time-variant factors include:

i) source-term characteristics
ii) frequency of natural events
iii) timing of events
iv) competency of engineered structures
v) population shifts

Each of these are discussed below.

FIGURE 2·4

POTENTIAL APPLICATION OF *DE MINIMIS* FOR DOSE TRUNCATION

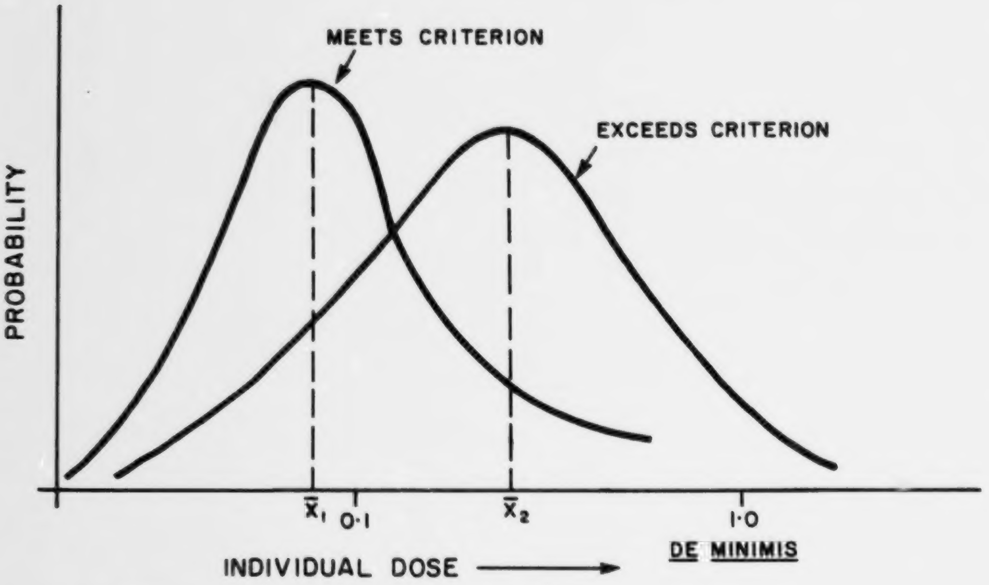

A) 1/10 *DE MINIMIS* CRITERION

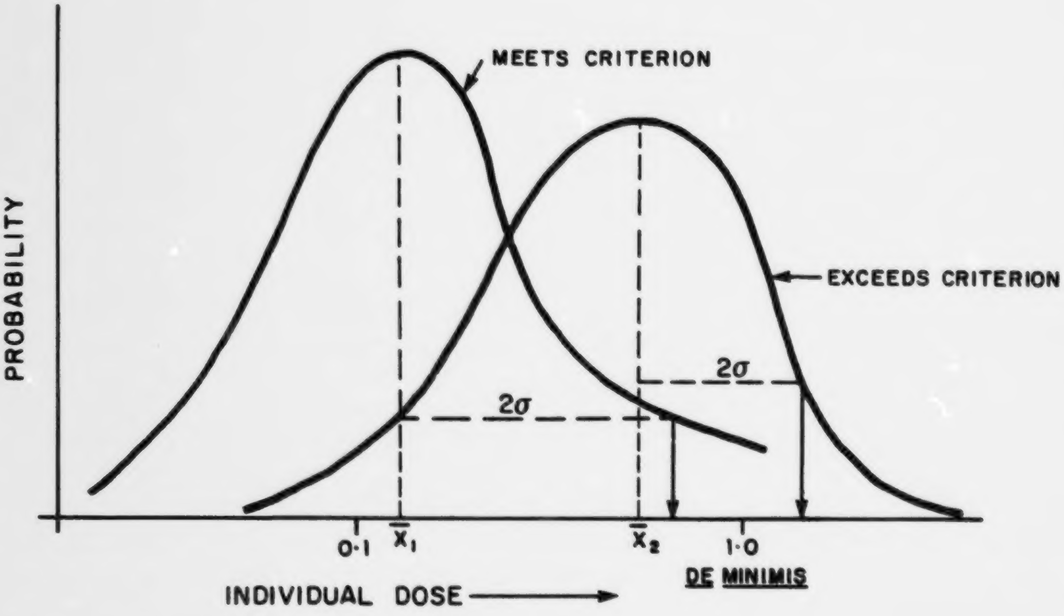

B) 95% CONFIDENCE CRITERION

2.3.1 Source-Term Characteristics

In the previous risk study (SENES 1985a), it was assumed that, in the absence of episodic events, the tailings did not change with time. However, in reality we know that the physical and chemical characteristics of a uranium tailings area will undoubtedly change with time. Many of these physical and chemical changes can be predicted and mathematically modelled. A previous study (SENES 1985b) predicted the time-varying behaviour of the source term and ultimately the radiation dose to a downstream receptor via a number of exposure pathways. The results of this study, illustrating the time-variant behaviour of dose to the receptor, are summarized on Figure 2.5. These latter, time dependent results are used to describe the null scenario (i.e. no episodic event scenario) in the present risk study.

The consequence (in this case, dose) obviously varies with time. Many of the pathways contributing to dose, namely dust, meat and vegetation are relatively constant with time. The doses from these pathways are all functions of the radionuclide concentrations at the tailings surface. The surface concentrations change very slowly, primarily as a result of radioactive decay.

The water-related pathways (water ingestion and fish) are dependent on far more complicated physical and chemical phenomenon. The increase of dose with time, and subsequent decrease, illustrated on Figure 2.5, is caused by the dissolution of radium associated with readily soluble gypsum in the tailings mass. Once this source of radium is depleted, the radium concentrations and hence the dose begin to decline.

Different physical and chemical phenomena will occur at different tailings areas, all resulting in somewhat different time-dependent behaviour.

2.3.2 Frequency of Occurence of Natural Events

In the Phase I risk study (SENES 1985a), several natural events (forest fire, drought and flood) were studied. The annual frequency of occurrence is

FIGURE 2·5

DOSE vs. TIME FOR NULL SCENARIO

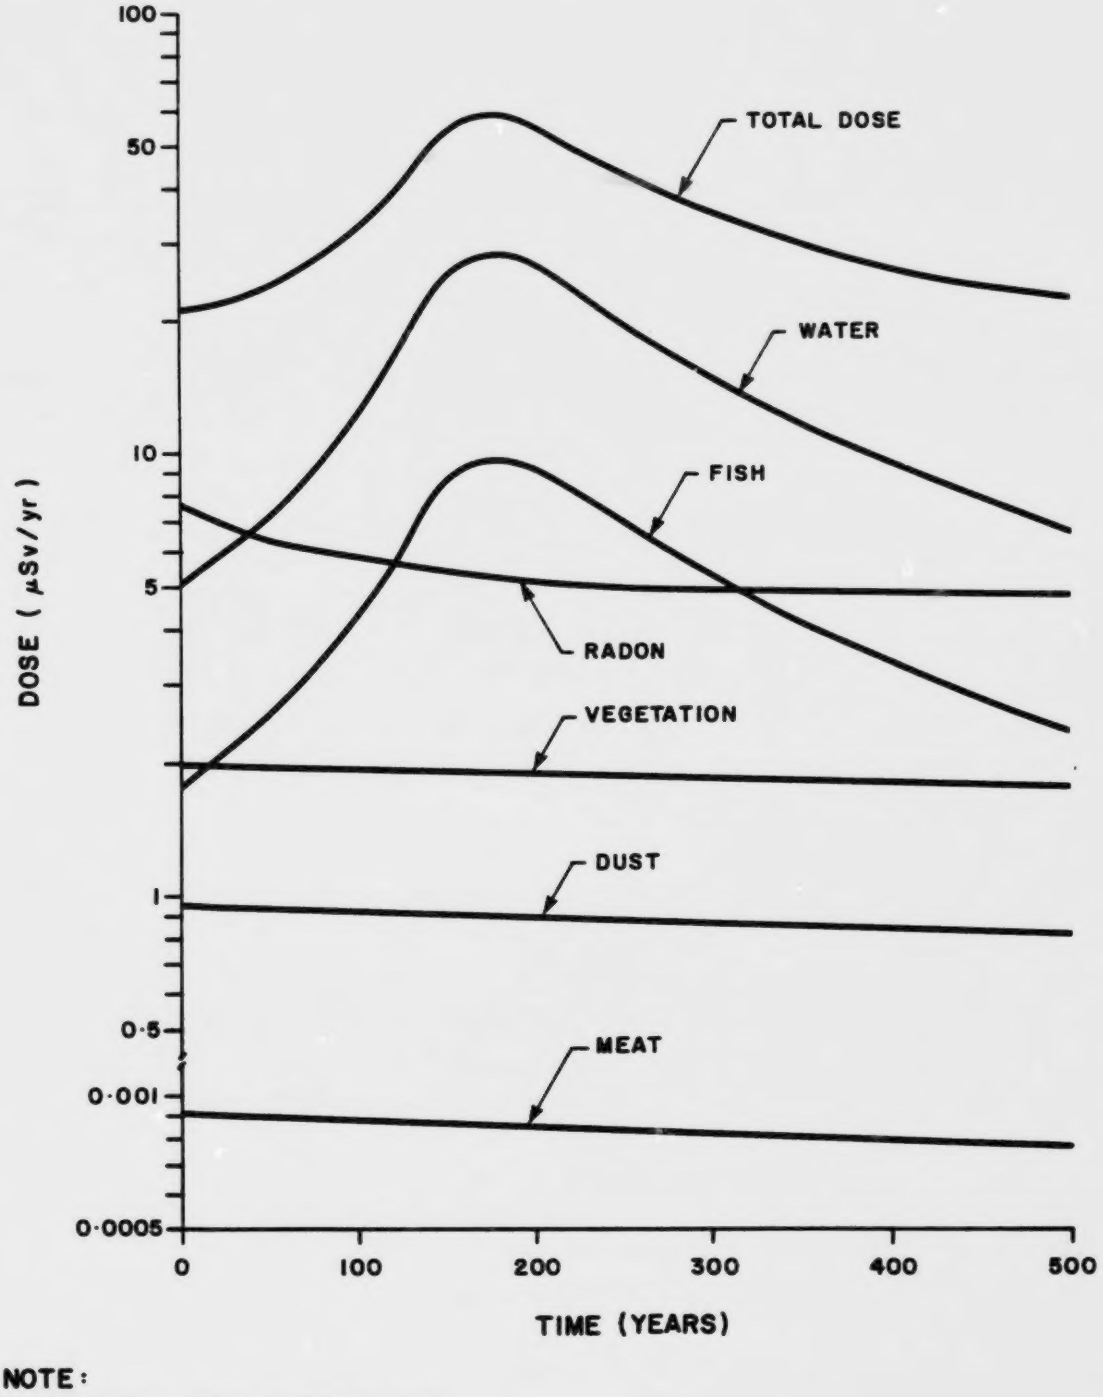

NOTE:

1. GEOMETRIC MEAN OF 500 TRIALS

SOURCE: AFTER SENES 1985, FIGURE 4·2

relatively well established for climate related events such as precipitation events, and major changes in climate are normally not considered for modelling periods of 1000 years or less. An excellent review of the frequency of occurrence of natural events is provided by McInnis (1985). Terms of record for any specific region are normally short; therefore, it is perhaps appropriate to acknowledge uncertainty in the annual frequency of occurrence. If risk assessment were extended into the 10,000 years range, then the effects of major climatic changes and aspects such as glaciation would need to be considered. Given the relative lack of knowledge on the frequencies of event occurrence, a range in values of one order of magnitude about the mean was assumed. It is also possible that the annual frequency of occurrence of such events may change in the future. However, whether and how such change would occur is unknown, adding yet another level of uncertainty about the occurence of future events.

2.3.3 Timing of Event Occurrence

Given that the source term characteristics are changing with time, it is apparent from Figure 2.5 that the magnitude of the consequences (e.g. annual dose to receptor) will also change with time. Thus, the consequences associated with the occurrence of an episodic event will be different depending on when the event occurs. This effect is illustrated schematically in Figure 2.6(a). We may therefore, choose to redefine an event (e.g. flood) in terms of when it occurs. Thus a flood in year 200 might be considered as a different event than a flood in year 800; in the sense that their consequences differ.

Alternatively, one may decide that the consequences over certain periods of time change so little that it does not matter in which particular year in that period the event actually occurs. Thus, one may decide to break the modelling period into "time windows", as shown in Figure 2.6(b). These time periods would likely be selected on the basis of the time-varying behaviour of the consequence(s) of interest; in the absence of any episodic events. This is described in more detail in a subsequent section.

FIGURE 2·6

EFFECTS OF TIME

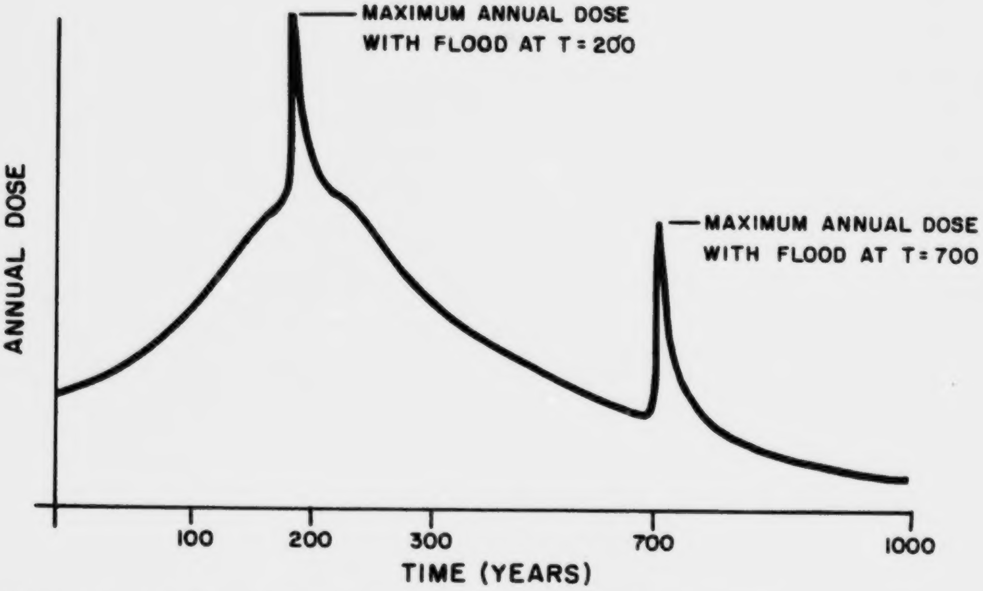

A) EFFECT OF TIME OF OCCURRENCE ON CONSEQUENCE

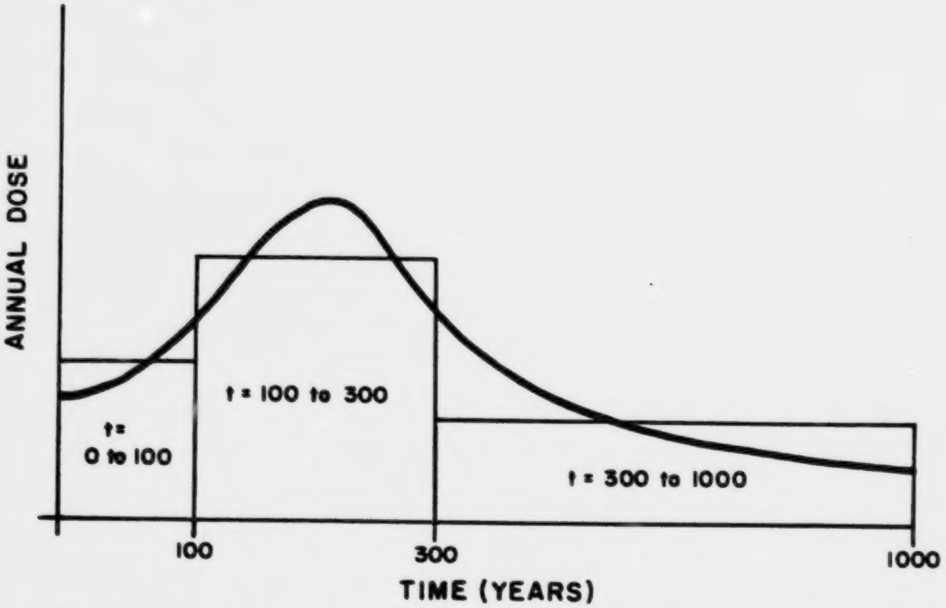

B) BREAKING THE MODELLING PERIOD INTO TIME WINDOWS

2.3.4 Competency of Engineered Structures

If structures such as dams, decants, etc. fail, it is likely either because of poor design/materials/construction techniques or as a result of long term degradation. Failures associated with design and construction faults are likely to occur in the 10-30 year operational life of the facility or not at all. Conversely, as structures degrade due to natural weathering from air and water erosion and chemical attack, the likelihood of failure increases. Thus, it seems reasonable to assume that design and foundation failures are not likely to occur, at least during the earlier portion of the simulation period, as the facilities have withstood 10-30 years of operations. As well, assuming institutional controls are in place for up to 200 years and that any damage from natural weathering could be identified and repaired if necessary, failures during this period are less probable. After the institutional control period ends, the chance of failure may increase with time.

2.3.5 Population Shifts

Major changes in population distribution can be expected over a 1000 year modelling period. Examining existing uranium mining areas gives some insight into the potential range of population changes. In the Uranium City area in northern Saskatchewan, the regional population varied from a small native population prior to 1949, to a major population center of more than 5000, to 3000 by 1980 and then to less than 300 in 1985. Similarly the Elliot Lake area consisted of remote cabins prior to 1955. Subsequently the population increased to a peak of more than 25,000 in 1960, reduced to less than 6000 by 1970 and increased again to more than 20,000 by 1980. Although this example is for population variations associated with mining activities and are for the short term it does illustrate the difficulty in projecting future population.

Because of the wide variations which have occurred and that are likely to occur, it is virtually impossible to make reliable predictions of future populations and consequently cumulative population dose. As such, the validity of using probabilistic risk assessments which attempt to use cumulative population dose as a true reflection of total risk can be questioned. This tool might however be useful for comparing differences between management options.

2.4 Incorporation of Uncertainty

Uncertainty is inherent in the concept of risk. There is uncertainty regarding the events that may occur, uncertainty regarding their annual frequency of occurrence and uncertainty in their consequences as predicted by mathematical models. The uncertainty regarding the predicted consequences results from the form of the mathematical models as well as from the values of parameters required as input to the models. The latter uncertainty can be accommodated explicitly by evaluating models probabilistically (using distributions of input parameter values) rather than deterministically (using single input values).

The Phase I risk study (SENES 1985a) ignored the uncertainties regarding both event frequencies and event consequences. The mean values of distributions of predicted consequences were used in the analysis. The results of this previous study were summarized by what is termed here a "level 1 risk curve".

To help us frame the question we wish to answer in this study, let us examine these uncertainties in a different fashion. It is possible to estimate a distribution of consequence for each event and/or each scenario at any time of interest within the simulation period. The distribution of consequence is conditional in the sense that it depends on the choice of the models and parameters used to calculate the consequences of a particular event; for scenarios, it is also conditional on the occurrence of a particular set of events.

The risk curve represents the probability (likelihood) that the various scenarios and their associated consequences actually occur. In symbols, the risk curve can be represented by:

$$P(C>C^*) = \sum_i P(C>C^*|S_i).P(S_i)$$

where

$P(C>C^*)$ is the probability of a consequence larger than any particular value C^*

$P(C>C^*|S_i)$ is the probability of a consequence larger than C^*, given a scenario "i" (i.e. it is a conditional probability) and,

$P(S_i)$ is the probability (likelihood) that scenario "i" will occur within the simulation period.

In the Phase I analysis, the distribution of consequence for a particular scenario $P(C>C^*|S_i)$ was represented by the mean-consequence $\overline{C}_i$ and the probability with which a particular scenario would occur within the simulation period $P(S_i)$ was assumed to be known with certainty and to be the likelihood l_i. In this case the risk curve, referred to as a level 1 risk curve, simplifies to:

$$P(C>C^*) = \sum_i l_i\overline{C}_i$$

This is, in fact, the risk curve discussed in Section 1.3 and illustrated on Figure 1.3.

However, we are actually uncertain both as to how likely a particular scenario is, and, if it should occur, what its consequences are. Expressed another way, we are somewhat uncertain about the actual shape of the risk curve, namely P $(C>C^*)$. Indeed by taking different combinations of likelihood and consequence it is possible to generate a whole series of risk curves and to express an opinion as to which curve is most likely. Alternatively it is possible to make statements about how confident we are that the "true" risk curve would lie within certain bands. The mechanics of this task form the basis for much of the remainder of this report. Specifically, in this study, we adopt an approach to the problem suggested by Kaplan and Garrick (1981). The methodology is described in some detail later, via a series of worked examples.

Basically, the uncertainty regarding event frequencies and consequences is represented using probability distributions of the quantities of interest. These distributions are then combined to produce a large set of level 1 risk curves. The probabilistic arithmetic can be carried out in various ways; however, discrete probability distributions were used in this study for mechanical convenience. The large number of level 1 curves are subsequently summarized and displayed as level 2 risk curves. The format of such curves are illustrated in Figure 2.7.

The ordinate on Figure 2.7 indicates the probability of exceeding any given level of damage. In keeping with the "probability of frequency" format of Kaplan and Garrick (1981) the vertical axis of Figure 2.7 is labeled frequency although the interpretation is the same. The curves also indicate the confidence we have in that estimate. That is, they indicate our degree of belief in the estimates we are making regarding the probability (relative frequency) of exceeding various levels of damage (the abscissa). For example, Figure 2.7 indicates that we are 90% confident that the frequency of exceeding C^* is not greater than $F_2(C^*)$ and 10% confident that the probability of exceeding C^* is not greater then $F_1(C^*)$. The latter concept may be restated as being 90% confident that the frequency of exceeding C^* is no smaller than $F_1(C^*)$.

FIGURE 2·7

GENERALIZED RISK CURVE

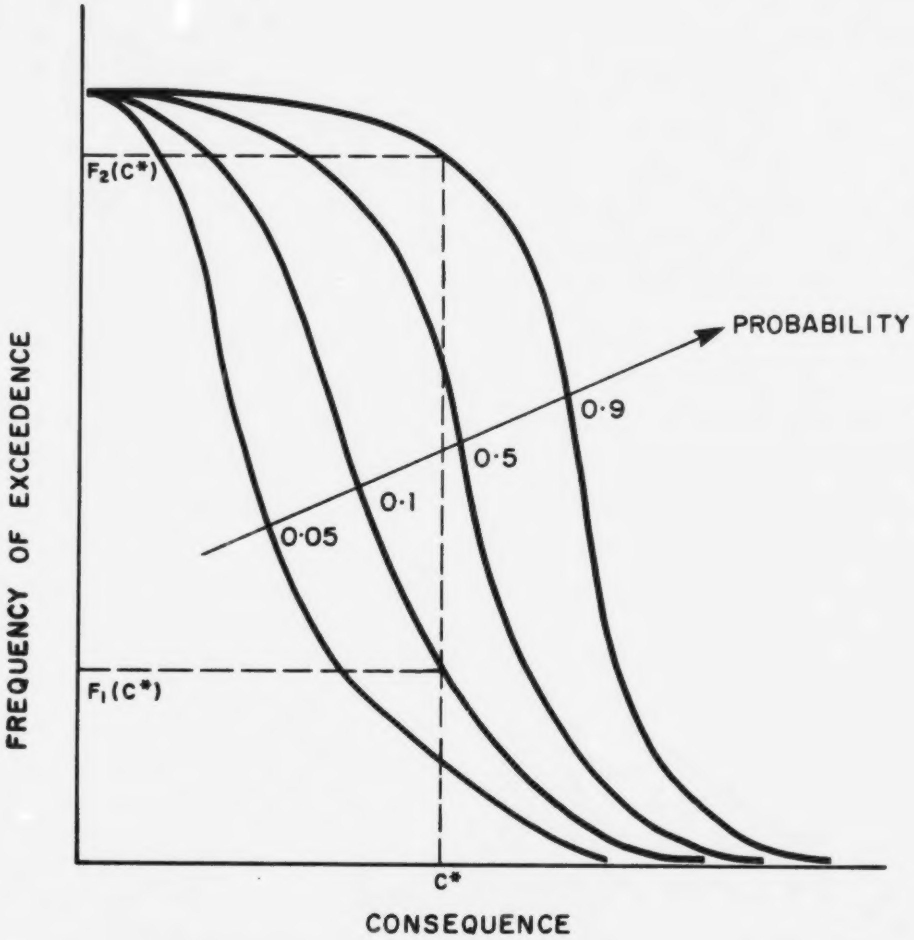

3.0 SITE DESCRIPTION AND DEFINITION OF EVENTS, SCENARIOS AND CONSEQUENCES

3.1 Description of the Reference Facility

To demonstrate the probabilistic risk assessment techniques a hypothetical reference tailings site was defined. The characteristics of this site, and the mathematical models used to describe its effects on the environment, were adopted from earlier NUTP projects (SENES 1985a,b.) The reference facility was defined to reflect the major environmental features and environmental transport pathways pertinent to uranium tailings management on the Canadian Precambrian Shield. The characteristics of this facility are illustrated in Figure 3.1.

The tailings disposal site was taken to occupy a natural depression formed within the upper reach of a valley, that is, a land/shallow-water type site. The tailings were considered to be contained behind a conventional dam constructed across the low point. The reference tailings site was defined to have a surface area of 50 ha and an average tailings depth of 10 m, for a total volume of 5×10^6 m^3. The natural depression containing the tailings was assumed to be 150 ha in area (including the tailings).

The reference tailings characteristics were assumed to represent those resulting from the acid leaching of low-grade pyritic uranium ore containing around 0.1% uranium. Tailings were assumed to contain 4 to 5% of the original uranium (i.e. between 95 and 96% of the uranium was extracted from the tailings during the milling process). Essentially all of the original radium in the ore typically remains with the tailings material. The radium concentration was selected to range between 5.0 and 15 Bq/g.

The effects of the tailings mass on the environment and on the hypothetical human receptors were determined using a set of environmental pathways models. Radionuclides modelled included radon gas release from the tailings surface, radium dispersal with windblown particulates and transport of radium dissolved in surface water. The exposure pathways considered are shown schematically in

FIGURE 3·1

SCHEMATIC ILLUSTRATION OF WATERSHED SETTING FOR THE REFERENCE TAILINGS SITE

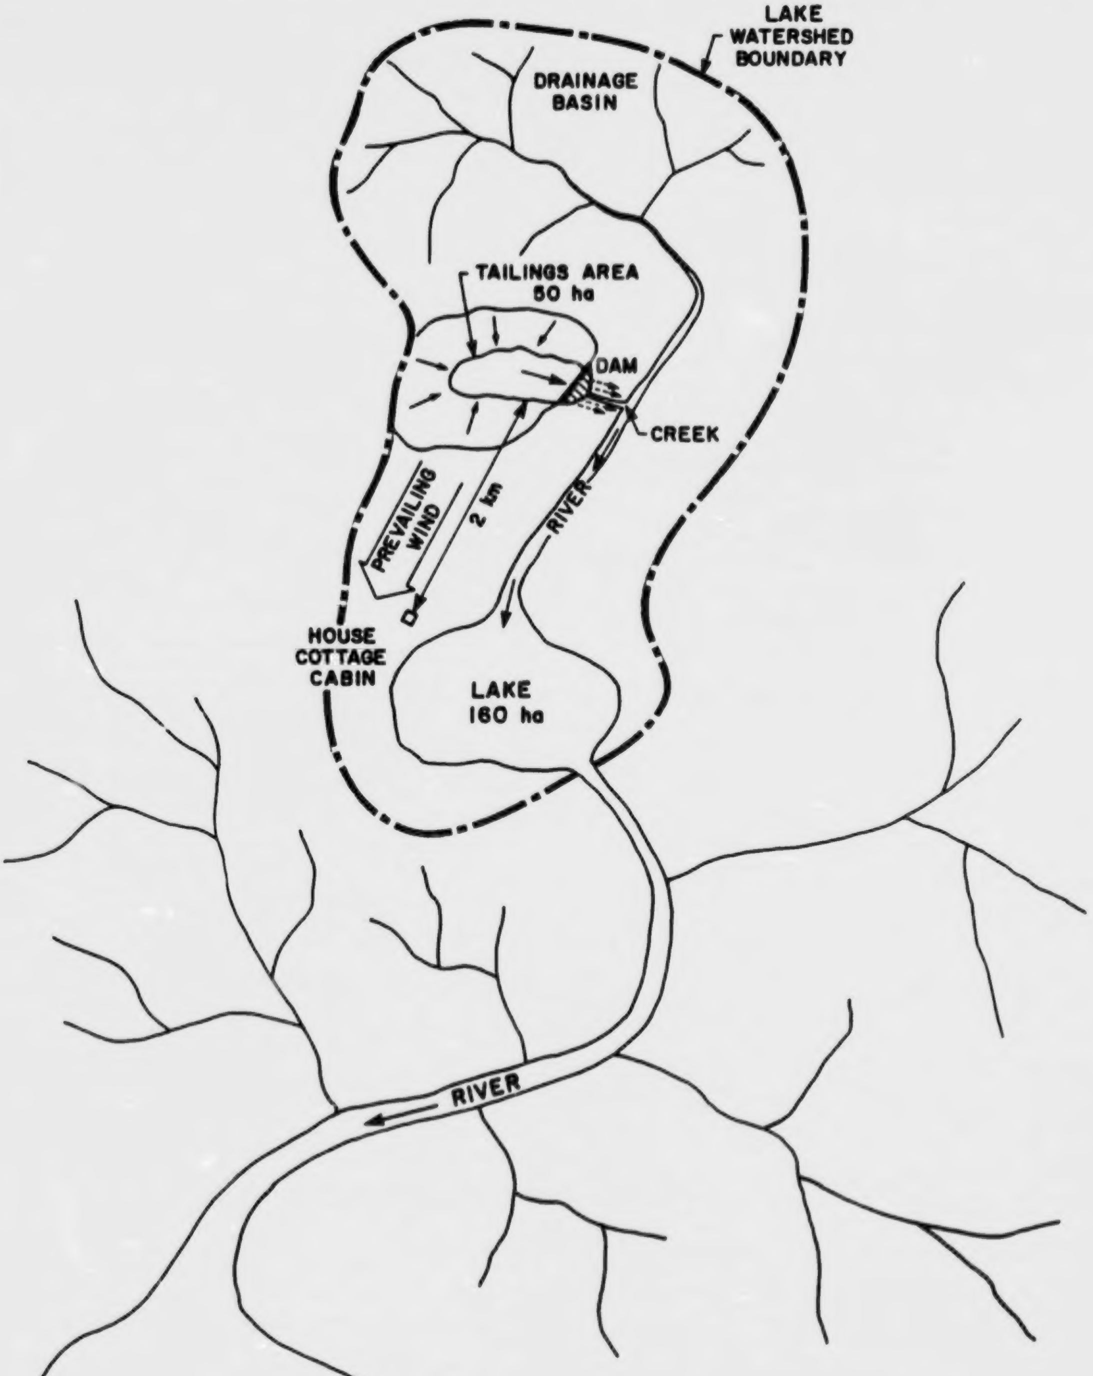
LAKE WATERSHED BOUNDARY
DRAINAGE BASIN
TAILINGS AREA 50 ha
DAM
CREEK
PREVAILING WIND
2 km
RIVER
HOUSE COTTAGE CABIN
LAKE 160 ha
RIVER

Figure 3.2 and described in some detail elsewhere (SENES, 1985a,b).

The number of pathways were limited to make the exercise tractable. However, it is worth noting that our experience suggests that even this restricted set of pathways illustrates all the salient features of pathways modelling and indeed is likely to account for the majority of the potential radiological health effects.

3.2 Events Considered

The basic "events" considered in this study are listed in Table 3.1 along with a nominal estimate of their annual frequency of occurrence and their probability of occurrence and non-occurrence over the thousand year modelling period. Such estimates are necessary in order to incorporate uncertainty in the risk analysis. As discussed in Section 1.3 the term probability often has the meaning of relative frequency of occurrence. This implies that this type of freqency (objective probability) can be determined, at least in principle, from a suitable experiment. However, with rare events, our experience base is small, and concepts such as best estimate and engineering judgements are used to establish subjective probabilities for the chance that an event will occur in any year. Thus, in this report the terms annual frequency of event occurrence or annual probability of event occurrence have been used more or less interchangeably. From here on however, we shall attempt, in talking about the annual occurrence of an event to use the term frequency or annual frequency. These numbers can then be combined in a suitable mathematical framework to calculate the overall chance that a particular combination of events will occur during the simulation period (ie. the scenario likelihood), Our uncertainty as to the annual frequency of occurrence is also discussed in Chapter 4.

FIGURE 3·2

EXPOSURE PATHWAYS ANALYZED

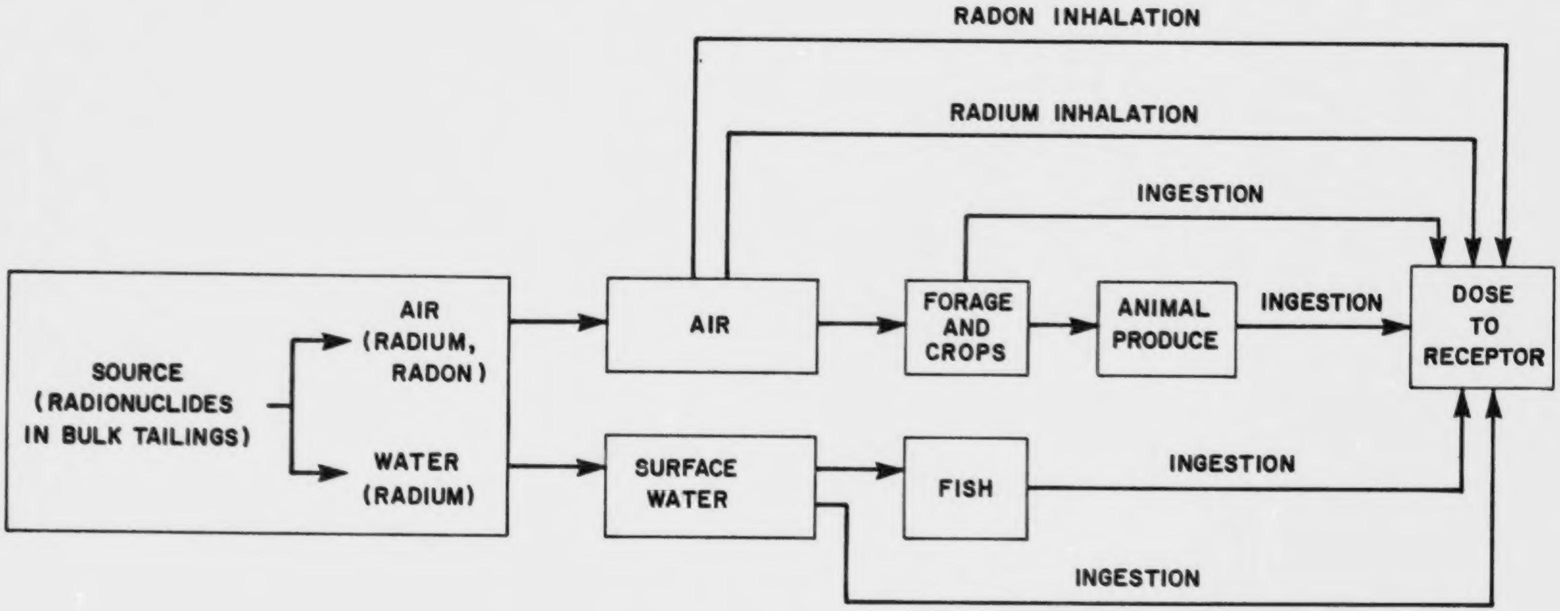
RADON INHALATION
RADIUM INHALATION
INGESTION
SOURCE (RADIONUCLIDES IN BULK TAILINGS)
AIR (RADIUM, RADON)
WATER (RADIUM)
AIR
FORAGE AND CROPS
ANIMAL PRODUCE
INGESTION
DOSE TO RECEPTOR
SURFACE WATER
FISH
INGESTION
INGESTION

Table 3.1

ASSUMED ANNUAL FREQUENCY OF EVENTS CONSIDERED

	Assumed Annual Frequency of Occurrence[1]	Frequency of Non-Occurrence in 1000 Years[2]	Frequency of One or More Occurrences in 1000 Years[3]
Natural Events			
Null[4]	0.998	NA	NA
Forest Fire	2×10^{-3}	0.135	0.865
Drought	1×10^{-4}	0.905	0.095
Flood	5×10^{-5}	0.951	0.049

Notes

1. Best estimate based on our current state of knowledge
2. $(1 - \text{annual frequency of occurrence})^{1000}$
3. Obtained by difference: (1 - frequency of non-occurrence)
4. Obtained by difference
5. Three figures are retained in the numbers shown in Table 3.1 for computation purposes. It is apparent that we do not actually know, or believe that the probabilities can be defined with such precision.

(i) The Null Event

The term "Null" event refers to the basic tailings scenario which, in the absence of one or more "unusual" events, is assumed to propagate through time in a more or less well-behaved fashion. In the Phase I risk study (SENES, 1985a), the null event was assumed to represent a constant steady-state situation. Specifically the null event consequences did not change with time. A subsequent study (SENES, 1985b) developed mathematical models which simulate the time-dependent behavior of the tailings.

(ii) Forest Fire

The effect of a forest fire is assumed to be a temporary reduction of vegetative cover with an attendant increase in windborne particulates and water erosion. For the present analysis this event is modelled as a reduction in vegetative cover from 90% (Null case) to 50% with natural revegetation taking place over a period of 10 years.

(iii) Drought

An extended period of severe drought is assumed to result in a substantial reduction in vegetative cover (from 90% to 10% cover). Following the drought there is a gradual recovery of vegetative cover over a 20 year period. During the period of reduced cover there will be increased wind and water erosion. During the drought the tailings moisture is reduced, leading to enhanced radon exhalation.

(iv) Precipitation/Containment Failure (Flood)

This event is actually the combination of two events; namely a major precipitation event (flood) leading to a failure of the tailings dam. A substantial fraction (10%) of the impounded tailings mass is assumed to be released to the lake. This results in an immediate increase in dissolved radium concentration in the lake; which is diluted relatively rapidly as a function of lake residence time. There will also be long-term leaching of radium from the redistributed tailings solids on the lake bottom. Radium flux will return to null event levels in approximately 100 years, as the lake bottom becomes covered with sediment. This event also results in increased wind and water erosion from the remaining, devegetated tailings mass. Revegetation, from the post-event 10% cover to the null event level (90% cover) is assumed to occur over a 20 year period.

3.3 Reference Receptor

The reference receptor is assumed to be an adult male who lives year-round in a cabin on the lake, as shown on Figure 3.1, some 2 km downstream of the reference tailings site. This individual is assumed to live off the land, deriving approximately 50 percent of his food supply from local sources: eating some vegetables and berries grown or harvested near the residence, eating some fish taken from the lake, and eating some caribou hunted in the immediate area. All drinking water is taken from the lake and consumed directly without special treatment for radium removal.

3.4 Complicating Factors Considered

As outlined in chapter two, there are many complicating factors which may be appropriate to consider in a risk analysis. Two of these, namely incorporating uncertainty and time dependence are of fundamental interest. Example applications which consider both uncertainty and time dependence presented in the following chapter were selected to demonstrate the flexibility of our approach to risk analysis. Although including each new factor further complicates the analysis, the basic approach and methodology remain the same. The factors included in each example are summarized in Table 3.2.

The first example examines the incorporation of uncertainty in both the frequency of event occurrence and in the event consequences. The consequence is taken to be the maximum annual dose to the reference receptor over the 1000

Table 3.2

FEATURES OF THE WORKED EXAMPLES

Example	Consequence	Time Variance	Uncertainty
1	Maximum Annual Dose to Receptor	No	Yes
2	Maximum Annual Dose to Receptor	Yes	Yes
3	Maximum Sulphate Concentration in Lake	Yes	Yes
4	Cumulative Dose	Yes	Yes

year modelling period. The example draws extensively on the results presented in the Phase I risk study (SENES, 1985a), and, in particular, assumes that the null event system is at constant, steady-state conditions. That is, no time variance in the null event dose is considered in the first example.

The second example also considers maximum annual dose to the receptor as the consequence of interest. In this example however, the null event dose is assumed to change with time (as shown in Figure 2.5). This example augments the results of the Phase I risk study (SENES 1985a) with results from another SENES study which modelled the time varying behaviour of uranium mill tailings (SENES, 1985b). As discussed previously, time dependency is incorporated by dividing the 1000 year modelling period into a series of "windows" in time. Similar events (eg. flood) occurring in separate "windows" are considered distinct and thus result in separate scenarios. As in the first example, uncertainty in both event frequency and consequence is incorporated in this example.

The third example duplicates many of the features of the second example, including a time varying null event and uncertainty in both event frequencies and consequences. The new aspect of this example is that a non-radiological consequence is considered. Namely, the maximum annual sulphate ion concentration ($[SO_4^{2-}]$) in the the 1000 year modelling period is used as the consequence of interest. Again, the results from a previous study (SENES 1985b) were used to estimate the consequences. This example indicates that the methodology is equally applicable to predicting chemical risks as it is to predicting radiological risks.

The first three examples all considered a "maximum annual" consequence as the damage of interest. Considering such consequences allows one to make many simplifying assumptions in performing the analysis. In particular, the occurrence of multiple events (e.g. 3 fires) in any scenario can be ignored, since the maximum single year consequence remains unchanged. However, when one considers consequences which accumulate over the 1000 year modelling period (e.g. cumulative 1000 year population dose) each occurrence of each event adds some increment of consequence to the scenario. Thus, the analysis becomes even more involved, and different simplifying assumptions are required to keep the analysis tractable. Example 4 illustrates this situation, using 1000 year cumulative dose as the consequence. The example also incorporates a time-varying null event and uncertainty in event frequencies and consequences.

4.0 DEMONSTRATION OF A METHODOLOGY FOR PROBABILISTIC RISK ASSESSEMENT

One of the primary objectives of this study was to develop a methodology for incorporating various uncertainties into our estimates of the risks associated with uranium mill tailings. The demonstration of the metholodology is by way of the four examples mentioned previously. The initial example is a step by step description of the process. The subsequent examples are presented in somewhat less detail.

Before outlining the methodology, we continue our review (begun in Section 1.3)of both the concept of risk adopted in this study and its graphical presentation. Following Kaplan and Garrick (1981) the estimates of risk (often displayed as "risk curves") which ignore uncertainty, are referred to as "level 1" estimates in the following text.

4.1 A Review of Relevant Concepts - Continued

In Section 1.3 we reviewed the basic idea of risk as a set of triplets and the concepts of event trees and risk curves. By way of specific example, let us consider an event tree generated in the Phase I Risk study (SENES, 1985a). The event tree shown in Figure 4.1 considers 3 events; fire, drought and flood; each with an associated annual frequency of occurrence and a consequence. The annual frequencies of event occurrence and the frequency with which an event could be expected to occur (or not to occur) during the full thousand year simulation period were taken from Table 3.1. These numbers were used to calculate the frequency with which the various scenarios (combinations of events) could be expected during the simulation period. If for example, we could perform an experiment where we could observe the behavior of say 10,000 tailings areas over a 1,000 year period, we would expect to observe Scenario 1 some 1160 times, Scenario 2 some 2320 times, etc.

The consequences were calculated using a set of environmental pathways models. One consequence, the maximum annual dose to the reference receptor, resulting from each event is summarized below:

FIGURE 4·1

EXAMPLE EVENT TREE

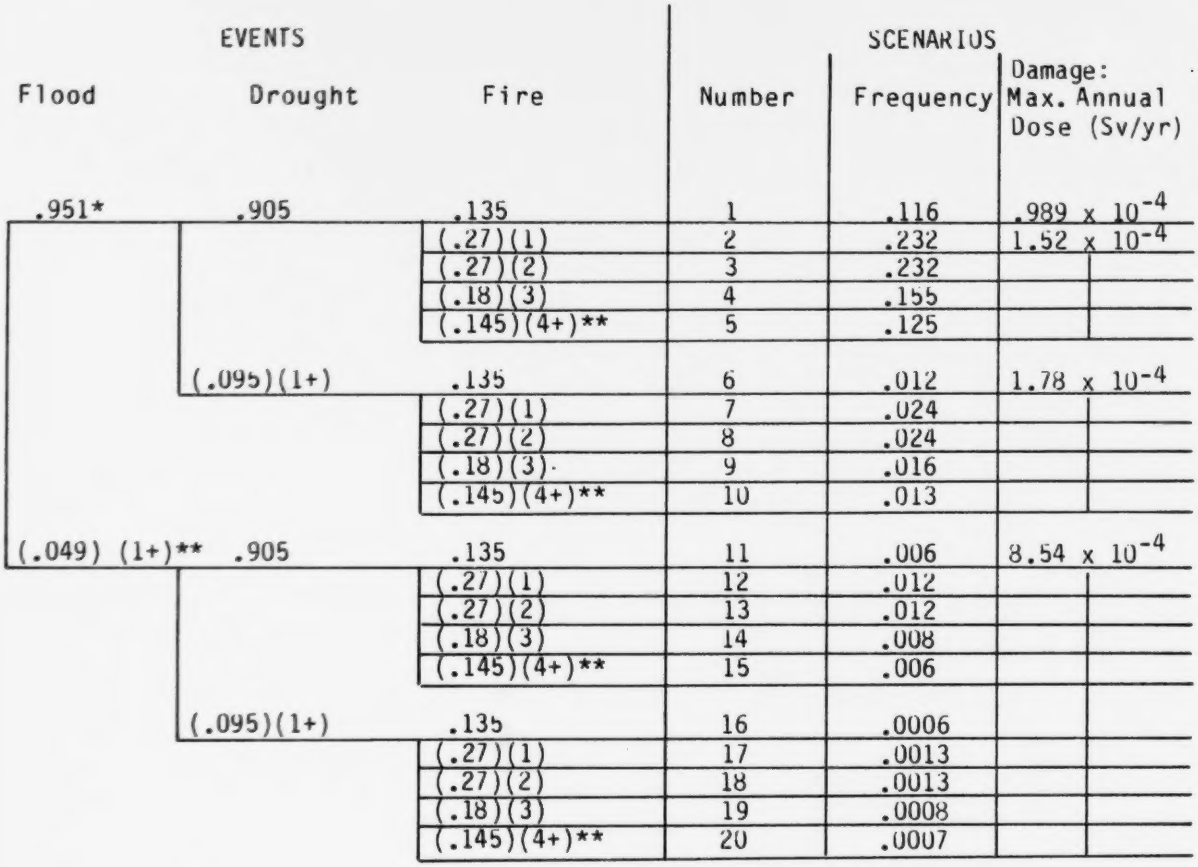

EVENTS			SCENARIOS		
Flood	Drought	Fire	Number	Frequency	Damage: Max. Annual Dose (Sv/yr)
.951*	.905	.135	1	.116	$.989 \times 10^{-4}$
		(.27)(1)	2	.232	1.52×10^{-4}
		(.27)(2)	3	.232	
		(.18)(3)	4	.155	
		(.145)(4+)**	5	.125	
	(.095)(1+)	.135	6	.012	1.78×10^{-4}
		(.27)(1)	7	.024	
		(.27)(2)	8	.024	
		(.18)(3)	9	.016	
		(.145)(4+)**	10	.013	
(.049) (1+)**	.905	.135	11	.006	8.54×10^{-4}
		(.27)(1)	12	.012	
		(.27)(2)	13	.012	
		(.18)(3)	14	.008	
		(.145)(4+)**	15	.006	
	(.095)(1+)	.135	16	.0006	
		(.27)(1)	17	.0013	
		(.27)(2)	18	.0013	
		(.18)(3)	19	.0008	
		(.145)(4+)**	20	.0007	

* Represents frequency of non-occurrence in 1000 years.

** Represents frequency of n or more occurrences in 1000 years.

Event	Maximum Annual Dose (Sv/yr)
null	0.989×10^{-4}
fire	1.52×10^{-4}
drought	1.78×10^{-4}
flood	8.54×10^{-4}

Note that the consequences are represented here and on Figure 4.1 by single numbers. These single numbers are in fact, the mean values of distributions of results. (It should be noted that the number of figures retained in the calculations do not indicate "significance" in the conventional sense. They are simply a numeric convenience.)

The "damage" considered in this example and noted on the event tree was the maximum single year dose to the reference receptor over the 1000 year modelling period. Thus, it was taken as the maximum annual dose to the reference receptor for the greatest consequence event in each scenario. For example, the maximum annual dose for scenario 7 (involving no flood, a drought and fire) is 1.78×10^{-4} Sv/yr; that resulting from a drought event.

Because of the particular damage index being considered, the event tree can be simplified considerably, to that shown on Figure 4.2. This simplication is possible since the damage associated with each scenario is dependent only on the highest consequence event in that scenario. That is, the number and occurrence of other events in the scenario can be neglected.

The risk can be presented graphically following the procedure outline by Kaplan and Garrick (1981) (SENES, 1985a). To construct a "risk curve", we determine the frequency that each dose (or greater) would occur over a 1000 year period. This is done simply by summing the scenario frequencies (or likelihoods) in the simplified event tree from highest to lowest damage. The resulting cumulative frequencies (which we interpret as cumulative (subjective) probabilities), are shown below and plotted as a Level 1 risk

FIGURE 4-2

SIMPLIFIED EVENT TREE

EVENTS			SCENARIOS		
Flood	Drought	Fire	Number	Frequency	Damage: Max. Annual Dose (Sv/yr)
.951	.905	.135	1	.116	0.989×10^{-4}
		.865	2	.745	1.52×10^{-4}
	.095		3	.090	1.78×10^{-4}
.049			4	.049	8.54×10^{-4}

curve in Figure 4.3.

Max. annual Dose (Sv/yr)	Frequency	Cumulative Frequency
8.54×10^{-4}	0.049	0.049
1.78×10^{-4}	0.090	0.139
1.52×10^{-4}	0.745	0.884
0.989×10^{-4}	0.116	1.0

4.2 Example 1: Incorporating Uncertainty

The previous analysis ignored the uncertainties which we know exist in both the frequency of event occurrence and in the consequences associated with each event. In this and subsequent sections, we will present a methodology for incorporating these uncertainties in our risk analysis. As indicated previously, the presentation will proceed by way of a worked example. The example used is an extension of that outlined in the previous section. The example will be presented in sufficient detail to allow all calculations to be replicated by the dedicated reader. It is anticipated that this level of detail, along with a summary of the methodology will provide the interested reader with sufficient information to repeat such an analysis.

The uncertainties in event frequencies and consequences can be represented as probability distributions of the appropriate quantities. Thus, symbolically, the actual event tree (for the previous example) might be constructed as shown in Figure 4.4. The uncertainties regarding the event frequencies are combined to produce the uncertainties regarding the scenario frequencies. The uncertainties regarding the event consequences can be determined via the probabilistic analysis of appropriate environmental pathways models.

Discrete probability distributions (DPD's) were used in subsequent calculations to carry out the probabilistic arithmetic. DPD arithmetic is described in detail elsewhere (Kaplan, 1981). This was done for convenience

FIGURE 4·3

LEVEL I RISK CURVE
MAXIMUM ANNUAL DOSE TO REFERENCE RECEPTOR

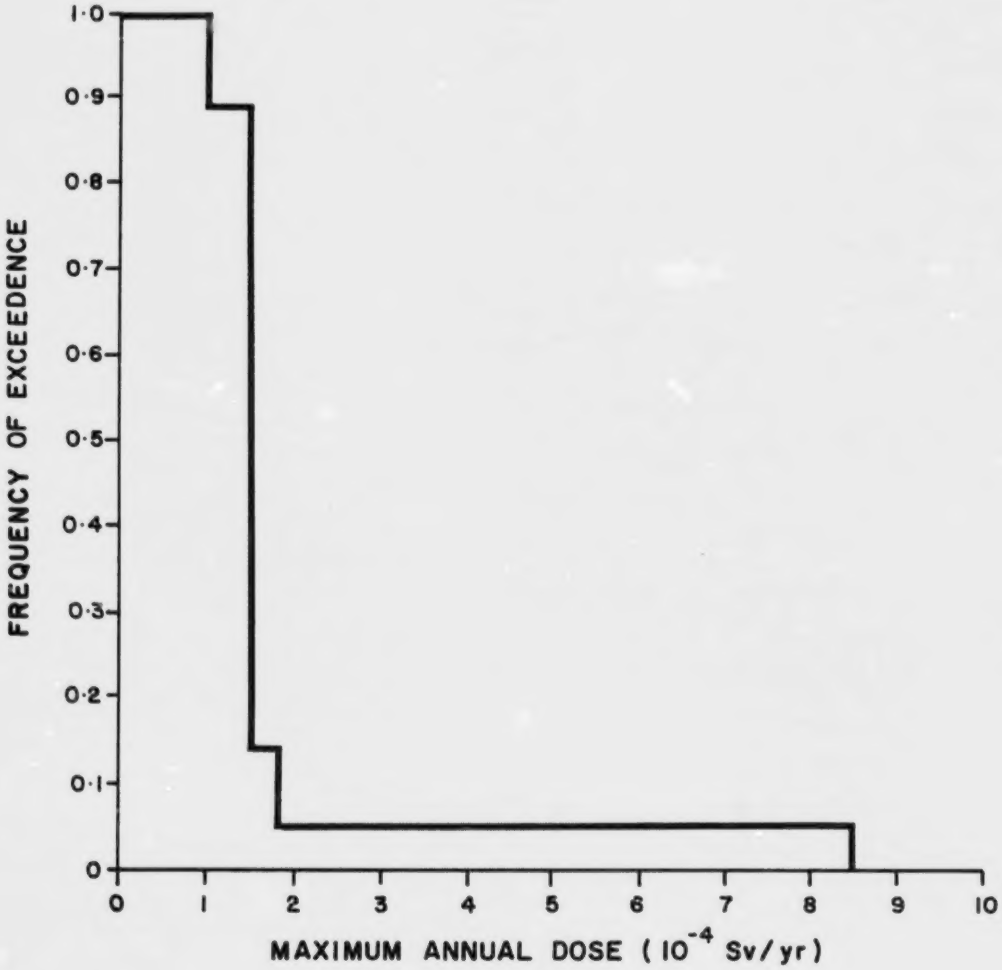

FIGURE 4·4

EVENT TREE - SYMBOLICALLY ILLUSTRATING UNCERTAINTY

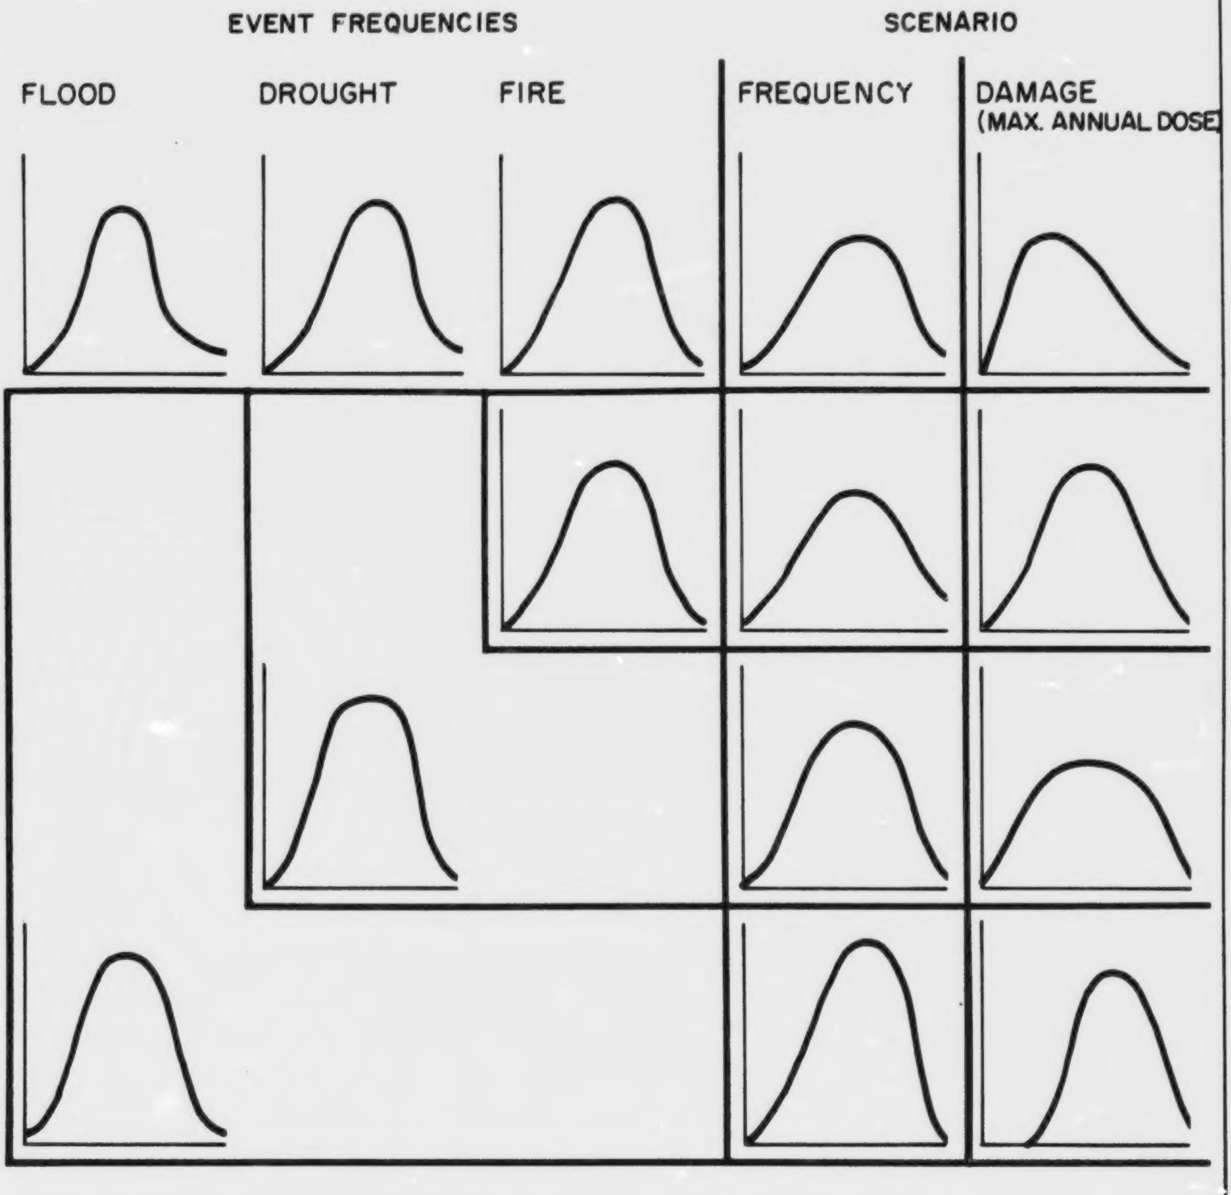
EVENT FREQUENCIES
SCENARIO
FLOOD
DROUGHT
FIRE
FREQUENCY
DAMAGE
(MAX. ANNUAL DOSE)

only. There is nothing intrinsic in the technique which requires DPD arithmetic.

4.2.1 Uncertainty in Scenario Frequencies

To incorporate uncertainty about the event frequencies, each was assumed to range over one order of magnitude. The values selected were factors of $\sqrt{10}$ and $1/\sqrt{10}$ above and below the previously specified value (Table 3.1). For purposes of illustration these ranges were specified using three point DPD's arbitrarily selected to have probabilities of 0.2, 0.6 and 0.2. (although any probability function could actually be used). The annual frequencies of occurrence were transformed into the frequencies of event occurrence and non-occurrence over a 1000 year period. These values, required for the calculation of scenario frequencies (likelihoods), are summarized in Table 4.1.

Thus, there are three possible frequencies at each branch point on the event tree. These event frequencies are treated as being completely independent; in the sense that learning more about the true frequency of forest fire would tell us nothing about the frequency of floods, etc. Thus, all possible event frequencies are combined with each other in determining the possible sets of scenario frequencies. The number of possible sets of scenario frequencies is thus:

$$\# \text{ of sets of scenario frequencies} = (\# \text{ of DPD points})^{\# \text{ events}}$$

In this case:

$$\# \text{ of sets of scenario frequencies} = 3^3 = 27$$

These sets of scenario frequencies are summarized in Table 4.2.

Note that the event frequencies shown on Table 4.2 are, in fact, the frequencies of one or more event occurrences over the 1000 years. To determine the scenario frequencies, either these numbers or the frequency of non-occurrence (i.e. 1-frequency of occurrence) must be used, as appropriate. For example, the frequency of scenario 1 (null event) in row one of Table 4.2

Table 4.1

UNCERTAINTY IN EVENT FREQUENCIES

Event	Annual Frequency	Probability of annual Frequency	Frequency of[(1)] non-occurrence in 1000 years	Frequency of[(2)] one or more occurrences in 1000 years
FIRE	6.3×10^{-4}	0.2	0.532	0.468
	2.0×10^{-3}	0.6	0.135	0.865
	6.3×10^{-3}	0.2	0.0018	0.999
DROUGHT	3.16×10^{-5}	0.2	0.969	0.031
	1.0×10^{-4}	0.6	0.905	0.095
	3.16×10^{-4}	0.2	0.726	0.274
FLOOD	1.58×10^{-5}	0.2	0.984	0.016
	5.0×10^{-5}	0.6	0.951	0.049
	1.58×10^{-4}	0.2	0.854	0.146

1) $(1\text{-annual frequency})^{1000}$

2) (1-frequency of non-occurrence)

Table 4.2

UNCERTAINTY REGARDING SCENARIO FREQUENCIES

Frequency of Event Occurrence			Resulting Sets of Scenario Frequencies (See Fig. 4.2)				Probability of Scenario Frequency
FLOOD	DROUGHT	FIRE	(1)NULL	(2)FIRE	(3)DROUGHT	(4)FLOOD	
0.016	0.031	0.468	0.507	0.446	0.0305	0.016	0.008
		0.865	0.129	0.825	0.0305	0.016	0.024
		0.999	0.001	0.952	0.0305	0.016	0.008
	0.095	0.468	0.474	0.417	0.093	0.016	0.024
		0.865	0.120	0.770	0.093	0.016	0.072
		0.999	0.001	0.890	0.093	0.016	0.024
	0.274	0.468	0.380	0.334	0.270	0.016	0.008
		0.865	0.096	0.618	0.270	0.016	0.024
		0.999	0.001	0.714	0.270	0.016	0.008
0.049	0.031	0.468	0.490	0.431	0.030	0.049	0.024
		0.865	0.124	0.797	0.030	0.049	0.072
		0.999	0.001	0.921	0.030	0.049	0.024
	0.095	0.468	0.458	0.403	0.090	0.049	0.072
		0.865	0.116	0.744	0.090	0.049	0.216
		0.999	0.001	0.860	0.090	0.049	0.072
	0.274	0.468	0.367	0.323	0.261	0.049	0.024
		0.865	0.093	0.597	0.261	0.049	0.072
		0.999	0.001	0.690	0.261	0.049	0.024
0.146	0.031	0.468	0.440	0.387	0.026	0.146	0.008
		0.865	0.112	0.716	0.026	0.146	0.024
		0.999	0.001	0.827	0.026	0.146	0.008
	0.095	0.468	0.411	0.362	0.031	0.146	0.024
		0.865	0.104	0.669	0.031	0.146	0.072
		0.999	0.001	0.772	0.031	0.146	0.024
	0.274	0.468	0.330	0.290	0.234	0.146	0.008
		0.865	0.084	0.536	0.234	0.146	0.024
		0.999	0.001	0.619	0.234	0.146	0.008

is:

$$(1-0.16) \times (1-.031) \times (1-.468) = .507$$

The frequency of scenario 2 (fire), in row one of Table 4.2 is:

$$(1-.016) \times (1-.031) \times (.468) = .446$$

Each set of scenario frequencies has associated with it a certain probability of occurrence. Consider, for instance, the probability of selecting the "low" frequency of occurrence for each of the three events, producing a set of "low" scenario frequencies and a risk curve with a "low" estimate of consequence frequencies. The probability of such an extreme set of scenario frequencies would be $(0.2) \times (0.2) \times (0.2) = 0.008$; as shown on the first row in Table 4.2.

4.2.2 Uncertainty in Scenario Consequence

Probabilistic evaluation of the models used in this study allows us to generate distributions of consequences (maximum annual dose in this example) for each of the events. These distributions represent our uncertainty regarding the consequence or damage associated with each event. As described previously, they are used to assign a damage to each scenario; completing the set of triplets.

Each distribution of maximum annual dose was summarized by a 5 point DPD with the discrete damage levels selected to represent the 10th, 30th, 50th, 70th and 90th percentiles of each distribution. Each point was assigned equal probability (0.2) in the DPD. This is in keeping with the interpretation of percentiles of a distribution. These damage levels are presented in Table 4.3.

In contrast to the independent treatment of event frequencies, the distributions for the damage indices are treated as totally dependent between

Table 4.3

UNCERTAINTY REGARDING SCENARIO CONSEQUENCES

Distribution of Consequence (max. annual dose; 10^{-4} Sv/yr)

	Percentiles				
Scenario	10th	30th	50th	70th	90th
(1)NULL	0.509	0.708	0.832	1.05	1.54
(2)FIRE	0.858	1.14	1.33	1.57	2.37
(3)DROUGHT	1.02	1.29	1.54	1.92	2.81
(4)FLOOD	5.09	6.10	7.24	8.74	12.6

scenarios. That is, only 5 "sets" of damage levels are considered to be possible, corresponding to the 5 percentiles mentioned earlier. The presentation of damage levels in Table 4.3 emphasizes this point. In assigning damages to each of the four scenarios in the event tree, entire columns, rather than single numbers are selected from Table 4.3.

The total dependence assumption can be rationalized in the following way. First, suppose that somehow we were to learn that for the flood scenario, the value of the maximum annual dose to the reference receptor was the 90th percentile value of the distribution. Since the same environmental models are used for the fire, drought and null events, we might be inclined to think that the true values for those scenarios would also be in the neighbourhood of the 90th percentile. This "total dependence" is a stronger assumption than we actually believe of course, but it tends to give conservative calculations and it tends to compensate for the total independence assumptions used with respect to the event frequencies.

4.2.3 Combining Uncertainties

At this point in the methodology, we have collected and tabulated all the information required to perform a "level 2" risk analysis. the uncertainty regarding event frequencies is summarized in Table 4.2, while the uncertainty regarding event consequences is presented in Table 4.3. These uncertainties must now be combined in some rigorous manner to produce level 2 risk curves.

First, we recall that each level 1 risk curve was comprised of a set of frequencies and consequences; one for each of the 4 scenarios in the simplified event tree. The information presented in Table 4.2 and 4.3 will allow the construction of a large number of Level 1 risk curves. In fact, each row in Table 4.2 will be combined with each of the 5 "sets" of damage levels in Table 4.3 to produce a total of:

$$27 \times 5 = 135 \text{ risk curves}$$

For illustrative purposes, several such risk curves are presented in Figure 4.5.

Curve 1 on the figure corresponds to frequencies from row 1 of Table 4.2 combined with damages from column 1 of Table 4.3. The probability associated with this curve is (0.2)(0.2)(0.2)(0.2)=0.0016. Curve 2 corresponds to frequencies from row 2 of Table 2 combined with damages from column 5 of Table 4.3; while curve 3 represents row 27 and column 3. In total, 135 such risk curves can be constructed from the information in Tables 4.2 and 4.3. Each of these curves has a probability associated with it. It is purely by chance that 2 of the 3 representative curves shown on Figure 4.5 have the same probabilities. If drawn on a common graph, these 135 curves would form a "mishmash" of lines, which we refer to as a "spaghetti of curves". This situation is illustrated symbolically on Figure 4.6. Having constructed this spaghetti of curves, the challenge now is to summarize or present the information in a more understandable, useable form. This step in the methodology is perhaps the most difficult to understand.

Conceptually, it involves looking at "slices" taken parallel to the vertical axis, through the set of curves. Such slices will intersect each of the 135 curves once, at some y (cumulative frequency) value. Each of the intersected curves will also have a probability associated with it. Thus, at each slice, representing a single consequence value, 135 pairs of (cumulative frequency, probability) values will be identified. This procedure is illustrated schematically in Figure 4.6.

As an example, consider a slice at an X (max. annual dose) value of 2.37×10^{-4}Sv/yr. It will intersect the three curves mentioned previously (Figure 4.5) to produce the following frequency, probability pairs:

FIGURE 4·5

SEVERAL REPRESENTATIVE RISK CURVES

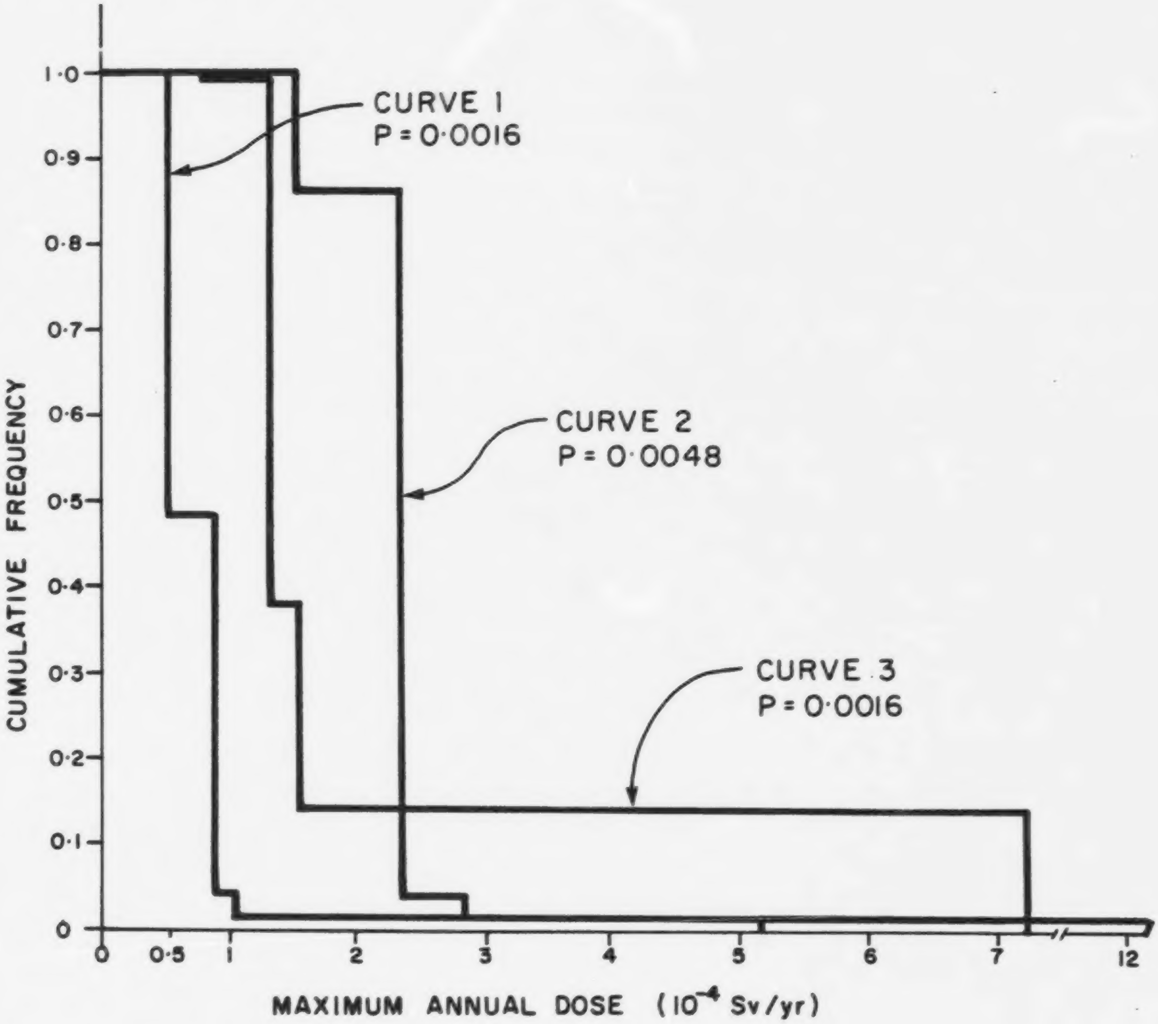

FIGURE 4·6

A "SPAGHETTI" OF RISK CURVES

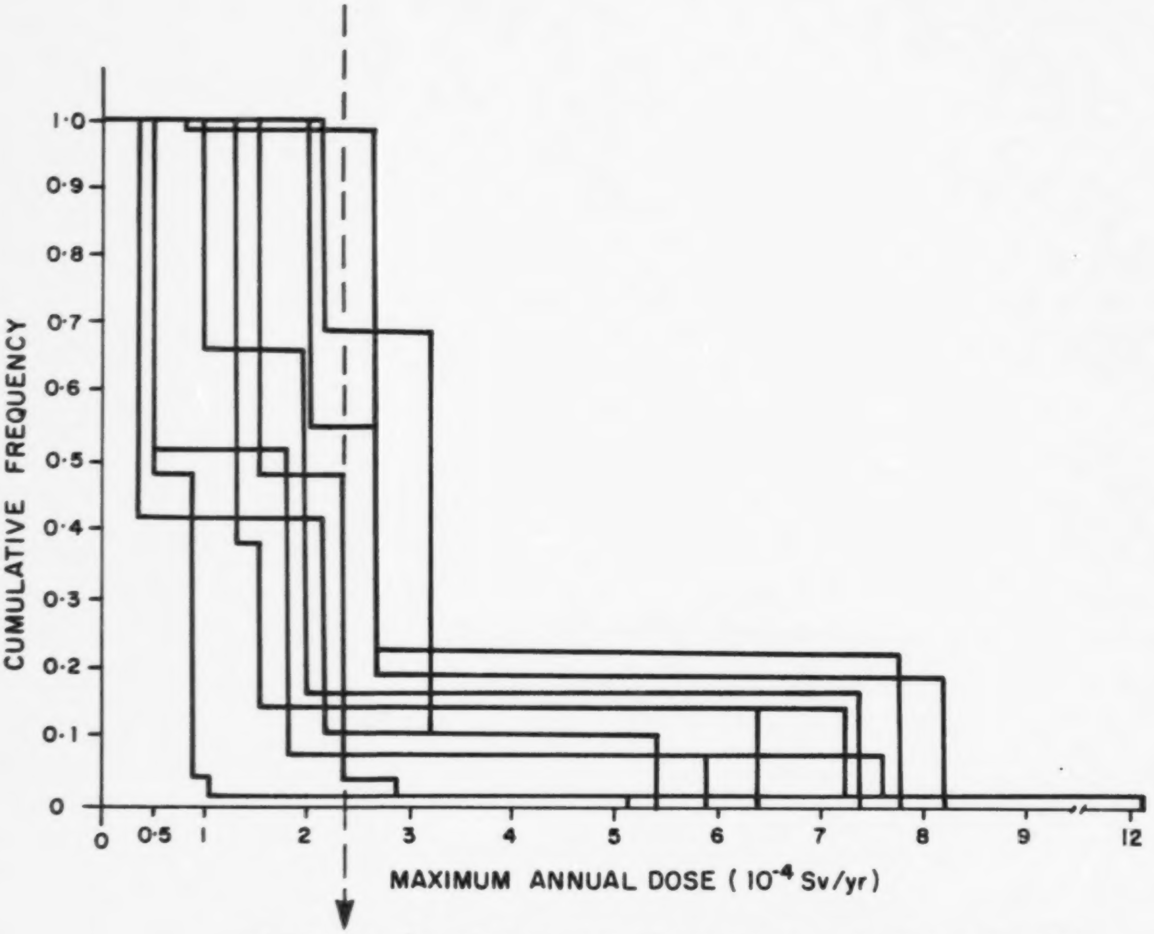

NOTE: AN "X VALUE SLICE" PRODUCES A SET OF (135) CUMULATIVE FREQUENCY, PROBABILITY PAIRS (SEE TABLE 4·4; FIGURE 4·7)

Curve	Cumulative Frequency	Probability
1	0.016	0.0016
2	0.871	0.0048
3	0.146	0.0016

Note that where the X slice corresponds exactly to a vertical step on the risk curve, the higher y (cumulative frequency) value is selected. In addition to these three, 132 additional pairs can also be determined; for each X slice.

Since 135 ordered pairs of numbers, at each of several X value slices, are not particularly informative, we wish to summarize them further. The following procedure applies to each X slice taken. It will be outlined below for only one slice (i.e. one set of 135 ordered pairs of (cumulative frequency, probability)) and must be repeated identically for each X slice taken.
To summarize the 135 ordered (cumulative frequency, probability) pairs, a cumulative density function is constructed; with cumulative frequency as the horizontal axis and cumulative probability along the vertical axis. This is accomplished by sorting the (frequency, probability) pairs in increasing order of cumulative frequency. Many of the cumulative frequency values will be the same, because of the discrete nature of the distributions used for event frequency. Cumulative probabilities are then calculated from the probabilities associated with each cumulative frequency. An example of this procedure for the X slice (2.37 x 10^{-4} Sv/yr) mentioned earlier is presented in Table 4.4. Once such a table is compiled the values may be used to construct a stepwise cumulative probability function, as shown on Figure 4.7.

The vertical axis of this cumulative curve represents the percentiles of the 135 cumulative frequency values. From the figure, the cumulative frequency values corresponding to each percentile (or cumulative probability) can be determined. Such values, for the deciles of the distribution presented in Figure 4.7 are tabulated in Table 4.5.

Table 4.4

SUMMARY OF CUMULATIVE FREQUENCY OF PROBABILITY FOR $X=2.37\times10^{-4}$

Cumulative Frequency	Number of Occurrences	Total Probability	Cumulative Probability
0.0157	36	0.160	0.160
0.0488	36	0.480	0.640
0.146	36	0.160	0.800
0.492	1	0.0016	0.802
0.509	1	0.0048	0.806
0.526	1	0.0048	0.811
0.542	1	0.0144	0.826
0.559	1	0.0016	0.827
0.589	1	0.0048	0.832
0.618	1	0.0016	0.834
0.631	1	0.0048	0.838
0.669	1	0.0016	0.840
0.871	1	0.0048	0.845
0.876	1	0.0144	0.859
0.880	1	0.0144	0.874
0.884	1	0.0432	0.917
0.888	1	0.0048	0.922
0.896	1	0.0144	0.936
0.903	1	0.0048	0.941
0.906	1	0.0144	0.955
0.916	1	0.0048	0.960
0.9983	1	0.0016	0.962
0.9983	1	0.0048	0.966
0.9984	1	0.0048	0.971
0.9985	1	0.0144	0.986
0.9985	1	0.0016	0.987
0.9986	1	0.0048	0.992
0.9987	1	0.0016	0.994
0.9988	1	0.0048	0.998
0.9989	1	0.0016	1.0

Table 4.5

DECILES OF THE CUMULATIVE PROBABILITY DISTRIBUTION FOR $X=2.37X10^{-4}$ Sv/yr

Cumulative Probability	Cumulative Frequency
0.1	0.016
0.2	0.049
0.3	0.049
0.4	0.049
0.5	0.049
0.6	0.049
0.7	0.146
0.8	0.146
0.9	0.884
1.0	0.9989

FIGURE 4·7

CUMULATIVE DISTRIBUTION OF PAIRS (CUMULATIVE FREQUENCY, PROBABILITY)

FOR X = $2{\cdot}37 \times 10^{-4}$ Sv/yr

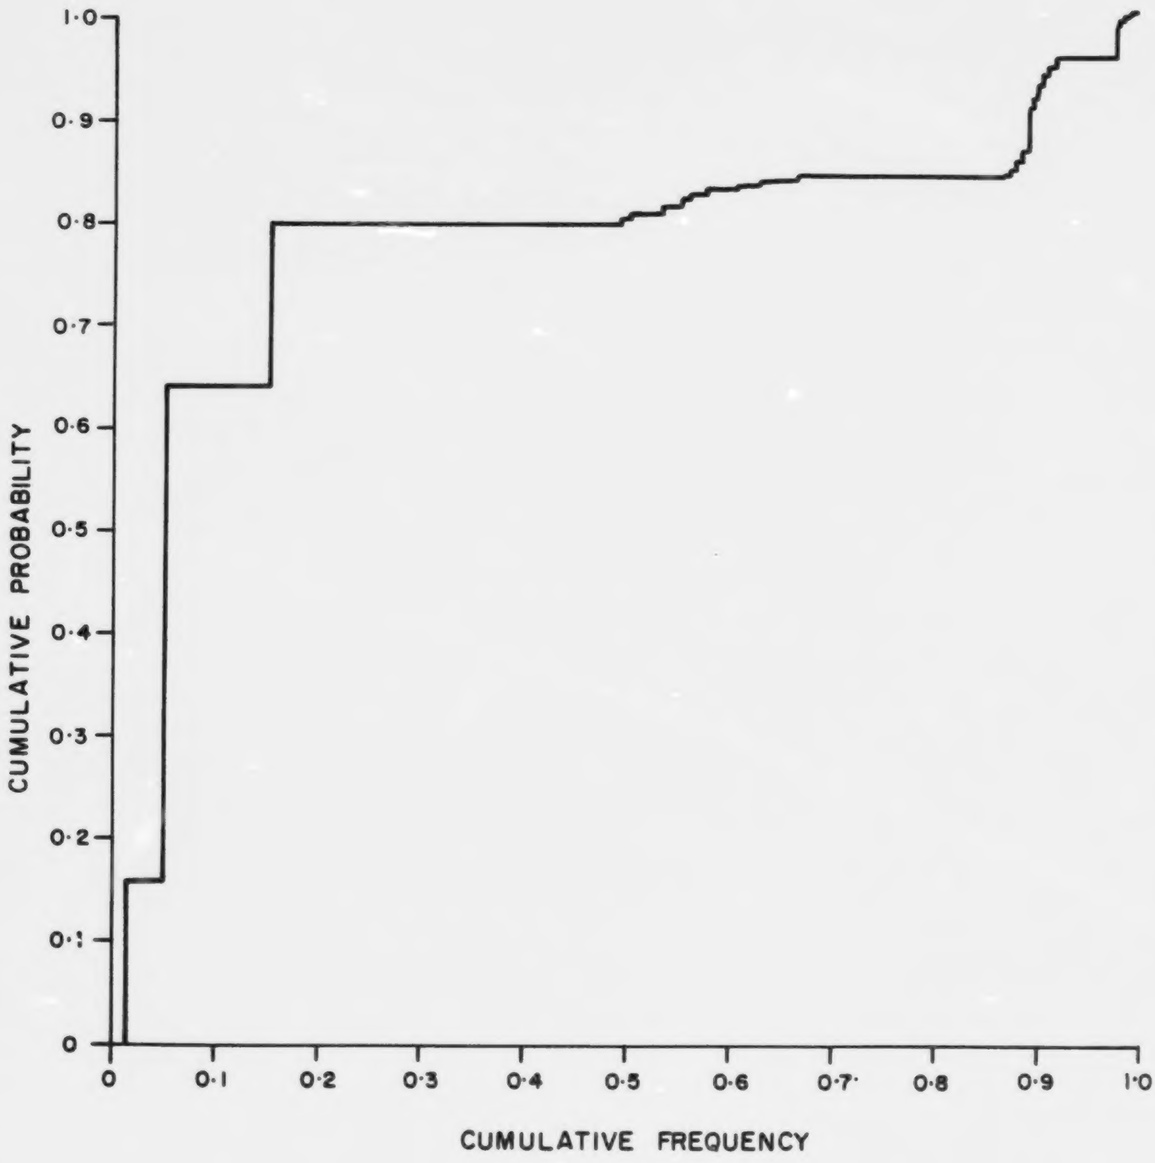

From this table, we can now return to our original coordinate system; with damage on the horizontal axis and cumulative frequency on the vertical axis. Selected deciles (the 10th, 50th and 90th percentiles) from Table 4.5 are plotted on Figure 4.8(a). Note that these values all correspond to a particular damage value (2.37×10^{-4} in this case). The points plotted represent the X (damage) and y (cumulative frequency) coordinates for each of three different curves (representing cumulative probabilities of 0.1, 0.5 and 0.9).

The entire procedure can be repeated at as many X values as desired. We repeated the procedure at each X value presented in Table 4.6. The 0.1, 0.5 and 0.9 probability values for each X slice are tabulated in Table 4.6 and displayed in Figure 4.8b. Note that, although the tabulated values rigorously define step-wise functions, the curves are plotted as smoothed curves on Figure 4.8b. We refer to the curves on the figure as a "family of risk curves".

To interpret these level 2 curves we refer to the underlying thought experiment (Kaplan and Garrick, 1981) in which we imagine building many millions of tailings piles and operating each for its 1000 year lifetime. We now ask the question, for example: "In what fraction of these lifetimes does a maximum annual dose of greater than 3×10^{-4} Sv/yr occur"

The answer to this question given by Figure 4.8b is: "We are 90 percent confident that this fraction is not bigger than 0.18. We are 50% confident that it is no bigger than 0.06, and 90% confident that it is no smaller than 0.03.

4.2.4 Summary of Methodology

Because of the large amount of information and large number of calculations inherent in the methodology, it is impractical to perform the required calculations by hand. To actually perform such a level 2 risk analysis, it is important to mechanize the procedure. Basically, the procedure involves the

Table 4.6

CUMULATIVE FREQUENCIES FOR SELECTED PERCENTILES AND DAMAGE LEVELS

Damage (Sv/yr)	Cumulative Frequency at Indicated Probabilities 0.1	0.5	0.9
0.509E-04	1.000	1.000	1.000
0.708E-04	0.884	1.000	1.000
0.832E-04	0.876	1.000	1.000
0.858E-04	0.618	0.996	1.000
0.102E-03	0.139	0.906	1.000
0.105E-03	0.049	0.906	1.000
0.114E-03	0.049	0.884	1.000
0.129E-03	0.049	0.876	1.000
0.133E-03	0.049	0.876	1.000
0.154E-03	0.049	0.146	1.000
0.157E-03	0.016	0.146	0.906
0.192E-03	0.016	0.078	0.884
0.237E-03	0.016	0.049	0.884
0.281E-03	0.016	0.049	0.146
0.509E-03	0.016	0.049	0.146
0.610E-03	0.000	0.049	0.146
0.724E-03	0.000	0.016	0.146
0.874E-03	0.000	0.000	0.049
0.126E-02	0.000	0.000	0.049

FIGURE 4·8

CONSTRUCTION OF RISK CURVES

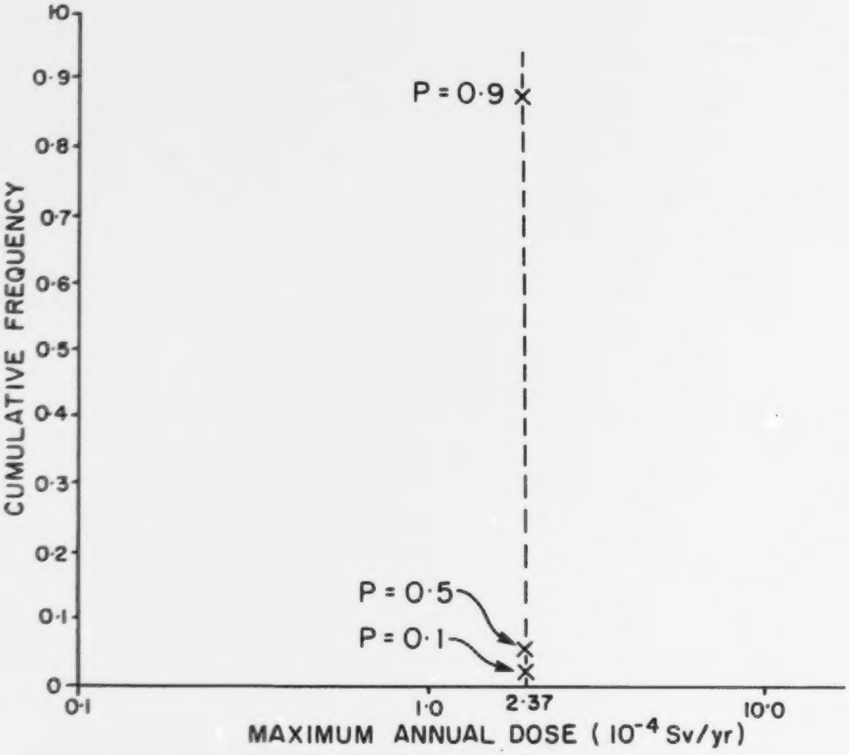

A) SELECTED PERCENTILES FOR A SINGLE X VALUE

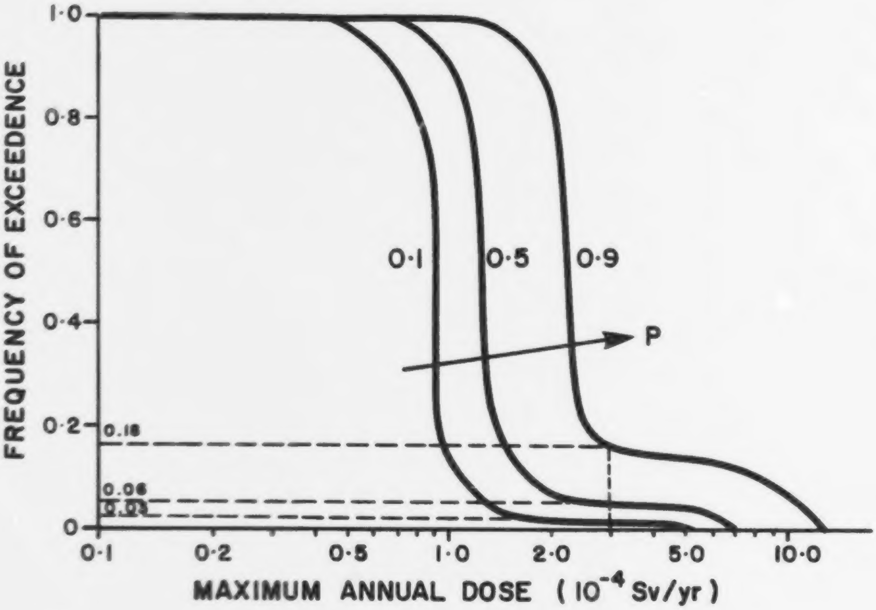

B) LEVEL 2 RISK CURVES

following steps.

(1) - construct event tree; showing relationships between all events and scenarios

(2) - identify DPD's for frequency of event occurrence and consequence; carry out the probabilistic arithmetic

(3) - construct risk curves representing all combinations of frequencies and consequences

(4) - take "X slices" through the spaghetti of risk curves. For each slice:
  - identify pairs of (cumulative frequency, probability); one pair for each risk curve
  - sort the pairs in order of increasing cumulative frequency
  - construct a stepwise cumulative distribution function (cdf); plotting cumulative frequency (horizontally) vs. cumulative probability (vertically)
  - extract from the stepwise cdf,the cumulative frequency values at selected probabilities
  - tabulate the cumulative frequency values for the selected probabilities at each selected x value.

(5) - these tabulated values represent a family of risk curves, and may be plotted as step-wise or smoothed curves.

4.3 Example 2: Incorporating Time Variance

The previous example assumed that the tailings, in the absence of episodic events would not change with time. That is, a steady-state null event was assumed. However, we know that, in reality, the physical and chemical characteristics of the tailings will change with time. A previous study (SENES 1985b) performed a probabilistic analysis of this time-varying behaviour. A summary of the results were presented earlier in Figure 2.5.

For convenience, this figure is reproduced here as Figure 4.9(a). As shown on the figure, the annual dose to the reference receptor increases during approximately the first 200 years and decreases thereafter.

It is obvious from the figure, and from previous discussion (Section 2.4) that event consequences will now be determined, in part, by the time of event occurrence. One could potentially try to include in the analysis, the possibility of each event occurring in each year. However, this would become an exceedingly complex problem. We have chosen instead, to divide the 1000 year modelling period into discrete "windows in time". As indicated on Figure 4.9(b), we have chosen to represent the continuous curve by three representative periods of time; thus creating a step-wise summary of the null event. Note that more and/or different "windows in time" could just as easily have been selected.

Given this approach to incorporating time variance, one can construct a new event tree, as shown in Figure 4.10. Note that since we are still considering maximum annual dose over the 1000 year as the consequence, the simplified event tree presented previously (Figure 4.2) forms the basis for the new event tree.

Time has been incorporated into the event tree, by expanding each scenario into three scenarios. Again, it should be noted that more and/or different windows in time could have been incorporated in exactly the same manner. Actual frequency values are not presented on the figure, since the frequencies of event occurrence will be considered uncertain in this example.

However, the frequency of each scenario is again calculated from the frequency of the event occurrences. The additional twist is the factor which accounts for the effects of time on scenario frequency. The effect of time is incorporated in each branch which indicates that an event has occurred. (i.e. the lower branches at each of the three event branch points.) At this point, we say that if an event has occurred, then it may have occurred in any one of the three time periods. Unless we are given reason to believe otherwise, we

FIGURE 4·9

INCORPORATING TIME VARIANCE

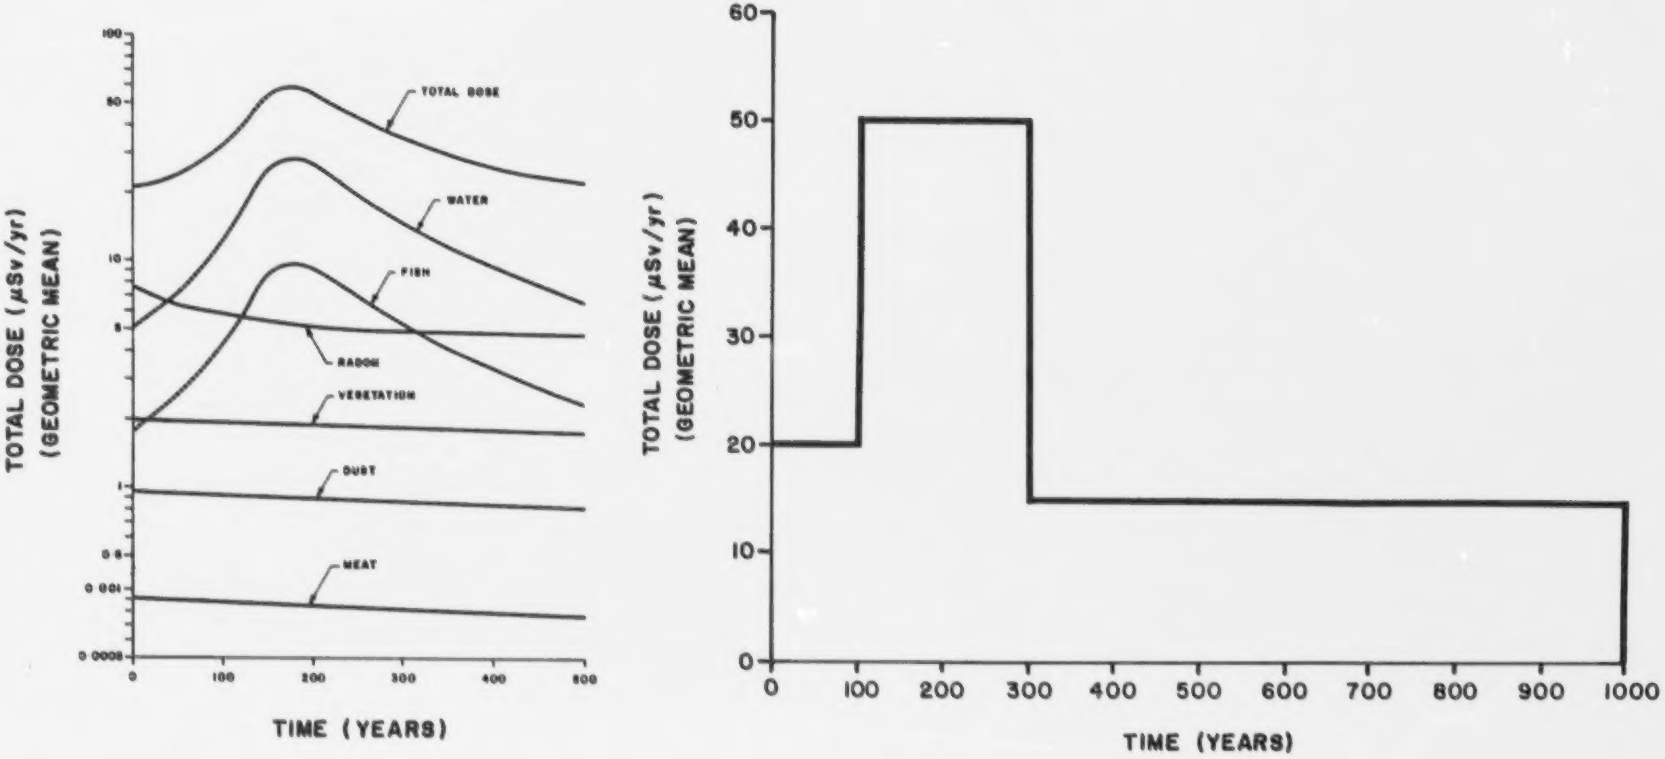

A) NULL EVENT: RECEPTOR DOSE vs. TIME

B) DISCRETE REPRESENTATION OF NULL EVENT TIME-VARIANCE

FIGURE 4·10

EVENT TREE:
INCORPORATING TIME-VARIANCE

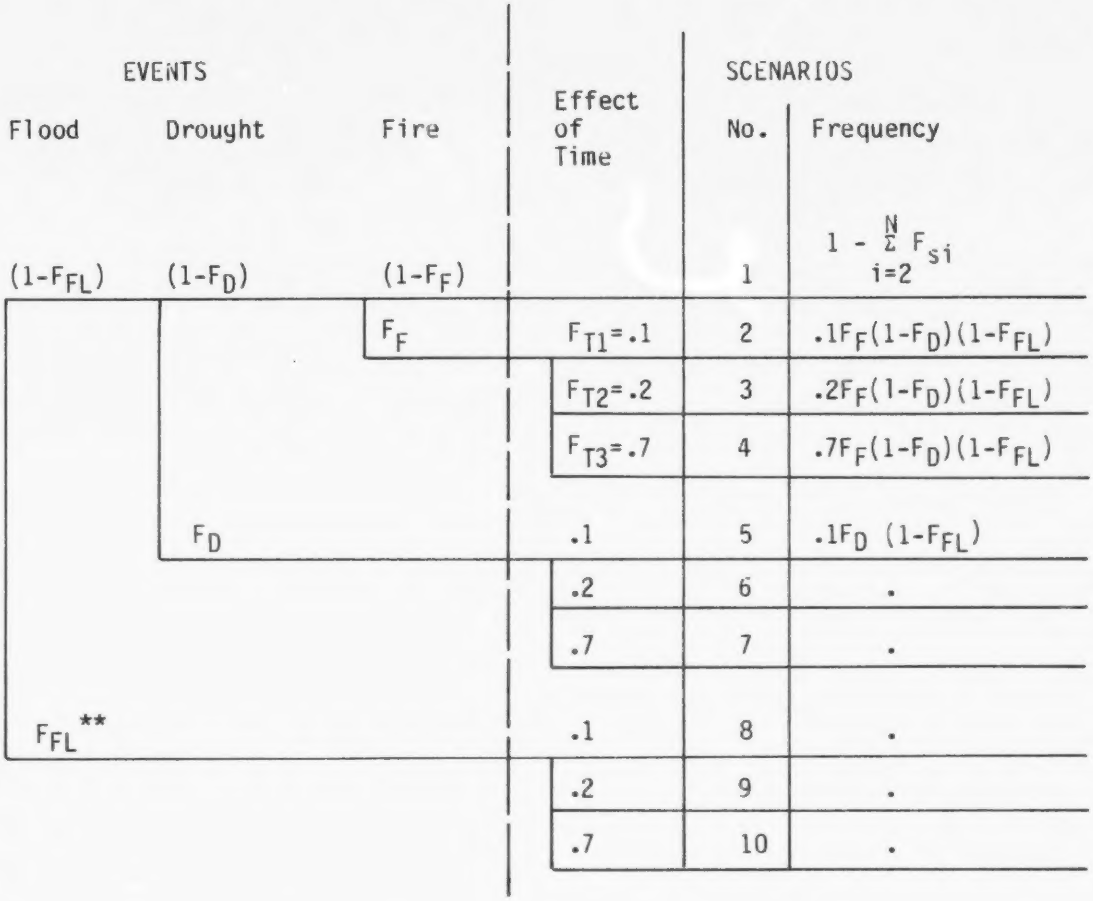
EVENTS
Flood
Drought
Fire
Effect of Time
SCENARIOS
No.
Frequency
$(1-F_{FL})$
$(1-F_D)$
$(1-F_F)$
1
$1-\sum_{i=2}^{N} F_{si}$
F_F
$F_{T1}=.1$
2
$.1F_F(1-F_D)(1-F_{FL})$
$F_{T2}=.2$
3
$.2F_F(1-F_D)(1-F_{FL})$
$F_{T3}=.7$
4
$.7F_F(1-F_D)(1-F_{FL})$
F_D
.1
5
$.1F_D\,(1-F_{FL})$
.2
6
.
.7
7
.
F_{FL}**
.1
8
.
.2
9
.
.7
10
.

* F_{si} = frequency of scenario i

** F_{FL} = frequency of one or more floods in 1000 years
F_D = frequency of one or more droughts in 1000 years
F_F = frequency of one or more fires in 1000 years
F_{T1} = frequency of an event ocurring in time period 1, (0-100 years) of the 1000 year simulation period etc.

assume that the events occur in the three time periods in direct relation to the length of each period. Hence, the chance of an event occurring between 0 and 100 years is 100/1000 or 0.1. Similarily, the chance that an event will occur in the 2nd and 3rd time periods are 0.2 and 0.7 respectively.

Should we have reason to believe that an event is more or less likely to occur in any time period, we would adjust the time factors accordingly. Only 2 of any 3 time factors may be adjusted independently; since the three (for any event) must always sum to 1.0. Hence if two factors are altered, the third has been implicitly determined.

It remains for us to add the damage index to the scenarios indicated in Figure 4.10. First, recall that the damage being considered in this example is the maximum single year dose to the reference receptor over the 1000 year period of each scenario. Since we are incorporating uncertainty regarding the consequences in this example, each scenario will have a distribution of values associated with it.

Consider first, the consequence associated with the null event. Each "window" or time period indicated on Figure 4.9(b) must be represented by a consequence, (i.e. annual dose). There are several possible approaches to selecting appropriate consequences for each time period. One approach would be to characterize each period by an average annual dose over that period. An alternative approach would be to characterize each period by the maximum annual dose observed in each time window. However, the approach selected here was to represent the consequences during each time period by those at the midpoint years. Thus, the probabilistic distributions of annual dose at t=50, t=200 and t=650 years were selected to represent the annual dose during their respective time periods. The distributions of values are summarized by their geometric mean values on Figure 4.9(b). It is clear from the figure that the maximum dose from the null event is that which occurs during the 2nd time period. This value (i.e. this distribution) of dose was used as the null event consequence in the subsequent analysis.

The damages associated with each of the events during each time period can be determined using methods outlined in the Phase I study (SENES, 1985a). This would consist of postulating each event occurring at each of 50, 200 and 650 years and examining the receptor doses during each of those years. We have amalgamated results from two previous studies (SENES, 1985a,b) to generate such damage distributions without actually performing computer simulations. The resulting distributions, summarized by their geometric means and geometric standard deviations are presented in Table 4.7.

It can be seen from Table 4.7 that several scenarios have consequences which are less than that of the null event. For instance, a fire in the first period results in a (geometric mean) dose of $.317 \times 10^{-4}$ sv/yr, which is less than the null event dose of $.495 \times 10^{-4}$ sv/yr. In fact, scenarios numbered 2, 4, 5, and 7 on Table 4.7 all have damages less than the null event. This indicates that in these scenarios (representing a 1000 year period) the maximum annual dose over the 1000 years arises NOT from the event occurrence but rather from the natural increase in null or "baseline" dose. Thus the damage (maximum annual dose over the 1000 year period) for these scenarios (2, 4, 5 and 7) should be $.495 \times 10^{-4}$ sv/yr.; that attributable to the null event.

Since this situation exists, the event tree shown in Figure 4.10 can be further simplified, by grouping all scenarios which are characterized by the null event dose. The simplified event tree is shown in Figure 4.11. The damages are summarized on the event tree by their geometric means and geometric standard deviations. The percentiles of the damage distributions are presented in more detail in Table 4.8.

As the event tree (Figure 4.11) indicates, we are now left with a total of six scenarios (including the null scenario). A frequency can be determined for each scenario, just as in the previous example. The frequency of the null scenario is simply; 1-(frequency of all other scenarios). The damages for each scenario are summarized as 5-point DPD's in Table 4.8. Thus, we are now ready to perform a level 2 risk analysis, using the procedure outlined in the previous section. The mechanical details of that procedure will not be

Table 4.7

SUMMARY OF TIME-VARYING DAMAGES[1]

Scenario No.	Description	Annual Receptor Dose(x 10^{-4} Sv/yr.) GM.	GSD.
1	Null	.495	2.11
2	Fire 1*	.317	1.66
3	Fire 2	.779	2.11
4	Fire 3	.234	1.87
5	Drought 1	.395	1.66
6	Drought 2	.924	2.11
7	Drought 3	.293	1.87
8	Flood 1	1.26	1.66
9	Flood 2	4.20	2.11
10	Flood 3	.902	1.87

* Refers to event in time period 1, 2 etc.

[1] The damage indicator used in the example was annual dose at the midpoint of the time window. As indicated in the text, other measures of annual dose could have been used with the methodology described.

Table 4.8

DISTRIBUTION OF CONSEQUENCES (MAXIMUM ANNUAL DOSE* (10^{-4}Sv/yr))

Scenario	Percentiles 10th	30th	50th	70th	90th
(1) Null	0.20	0.33	0.52	0.73	1.30
(3) Fire 2	0.32	0.54	0.76	1.17	2.00
(6) Drought 2	0.30	0.62	0.90	1.38	2.38
(8) Flood 1	0.68	0.98	1.27	1.65	2.39
(9) Flood 2	1.60	2.90	4.10	6.10	10.9
(10) Flood 3	0.40	0.65	0.90	1.25	1.98

* maximum annual dose at mid-point of each time window

FIGURE 4·11

SIMPLIFIED EVENT TREE: INCORPORATING TIME VARIANCE

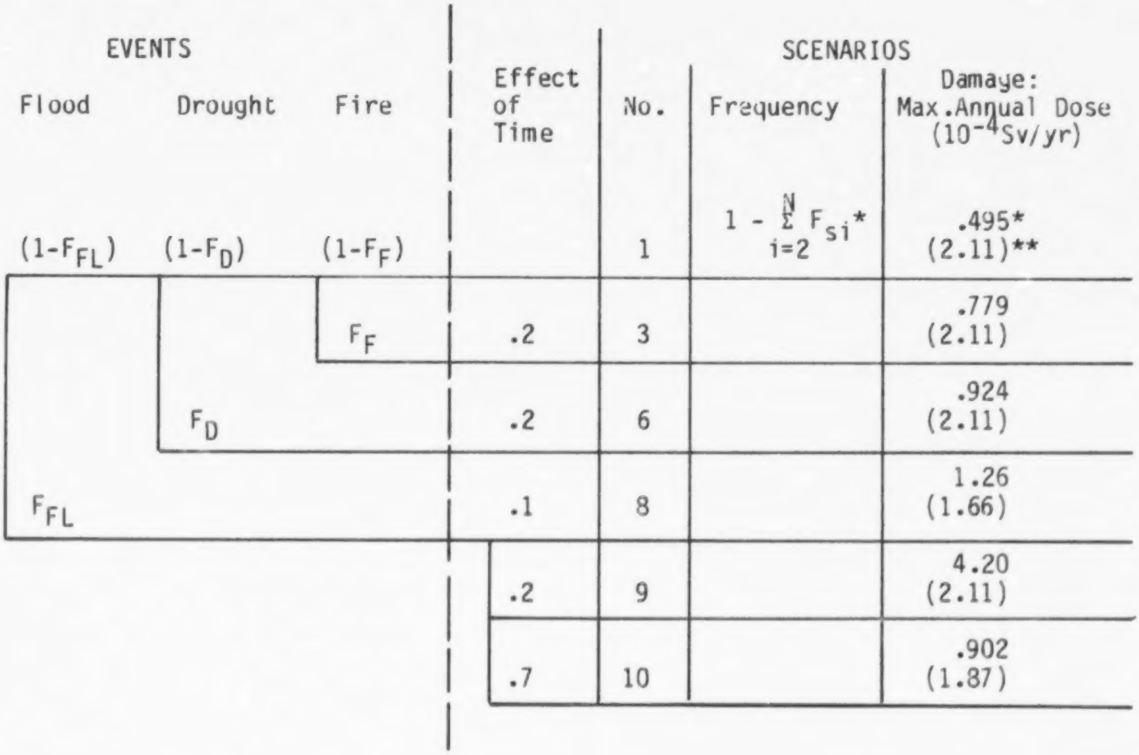
EVENTS
Flood
Drought
Fire
Effect of Time
SCENARIOS
No.
Frequency
Damage: Max. Annual Dose (10^{-4} Sv/yr)
$(1-F_{FL})$
$(1-F_D)$
$(1-F_F)$
1
$1 - \sum_{i=2}^{N} F_{si}$*
.495*
(2.11)**
F_F
.2
3
.779
(2.11)
F_D
.2
6
.924
(2.11)
F_{FL}
.1
8
1.26
(1.66)
.2
9
4.20
(2.11)
.7
10
.902
(1.87)

* Geometric mean
** Geometric standard deviation

repeated here. A spaghetti of:

$$3^3 \times 5 = 135$$

level 1 risk curves were generated, by again using 3-point DPD's for the event frequency distributions. These were summarized using the previous approach to produce the level 2 risk curves shown in Figure 4.12.

4.4. Example 3: A Non-Radiological Consequence

The approach to risk assessment being used in this study has been applied to many different problems. To demonstrate its flexibility, the risk arising from a non-radiological consequence is considered in this example. The consequence here is taken to be the maximum annual concentration of sulphate ion ($[SO_4^{2-}]$) over the 1000 year modelling period, in a lake just downstream of the tailings. Again, because a maximum, rather than a cumulative consequence is being considered, multiple event occurrences may be ignored. If we choose to divide the 1000 year modelling period into the same 3 "windows" as previously, the structure of the event tree for this example will be the same as that in the previous example (Figure 4.10).

Thus, the scenarios and their frequencies are defined as in the previous example. It remains for us to generate damage distributions for these scenarios with maximum sulphate concentration as the consequence of interest. Again, no explicit modelling was done. The damage distributions were generated using the results of 2 previous studies (SENES, 1985a,b).

The change in null event sulphate concentration with time is illustrated as a step function in Figure 4.13. Again, the distributions of sulphate concentrations at t=50, 200 and 650 years are summarized by their geometric means and geometric standard deviations. Since the sulphate concentration decreases from 50 to 200 to 650 years, the null event sulphate concentration is that calculated in the first time period. The distributions of sulphate concentration for each of the events, for each of the three time periods

FIGURE 4·12

EXAMPLE 2 : RISK CURVE

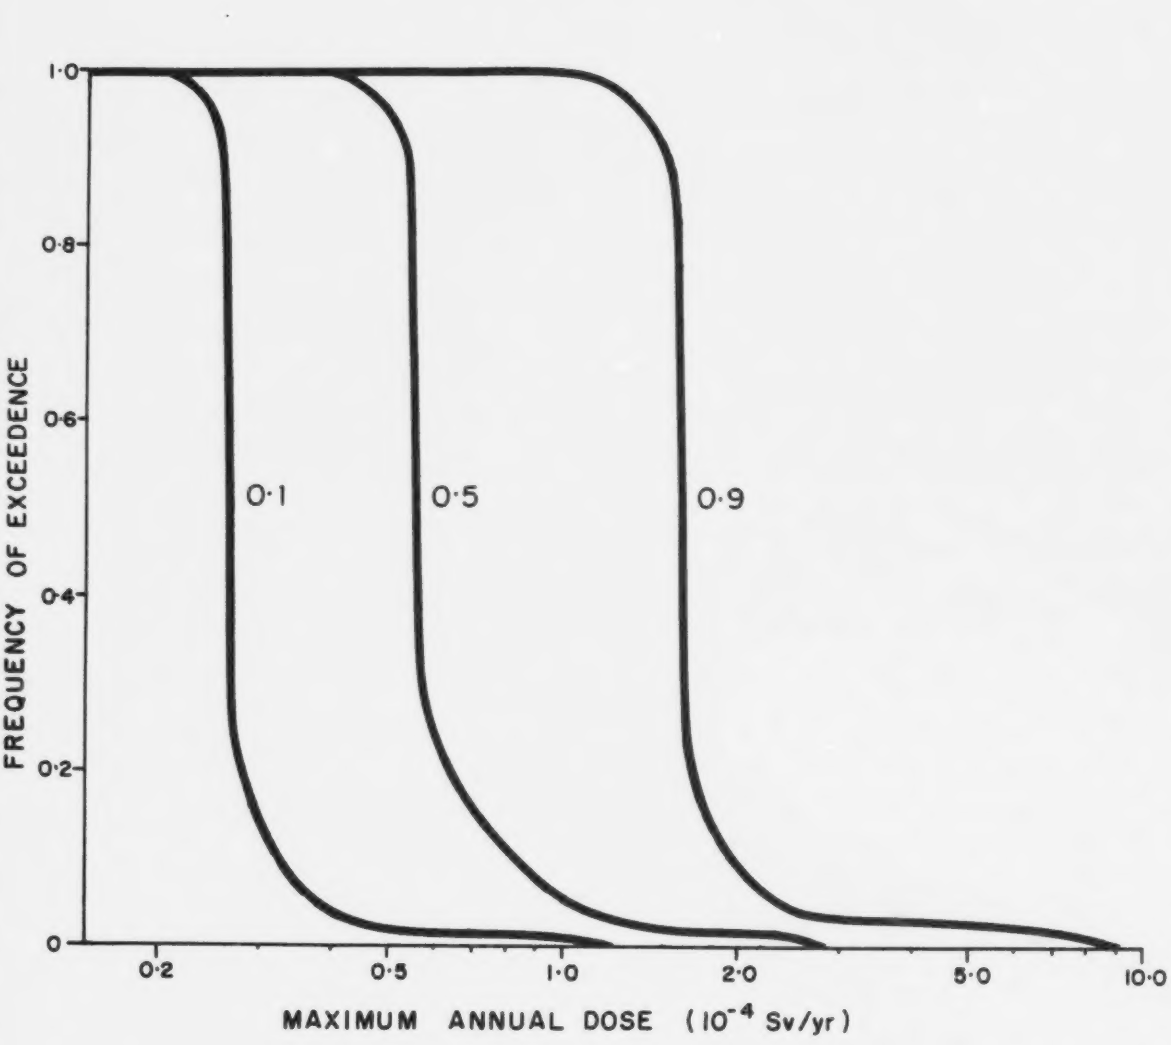

FIGURE 4·13

TIME VARYING SULPHATE CONCENTRATION

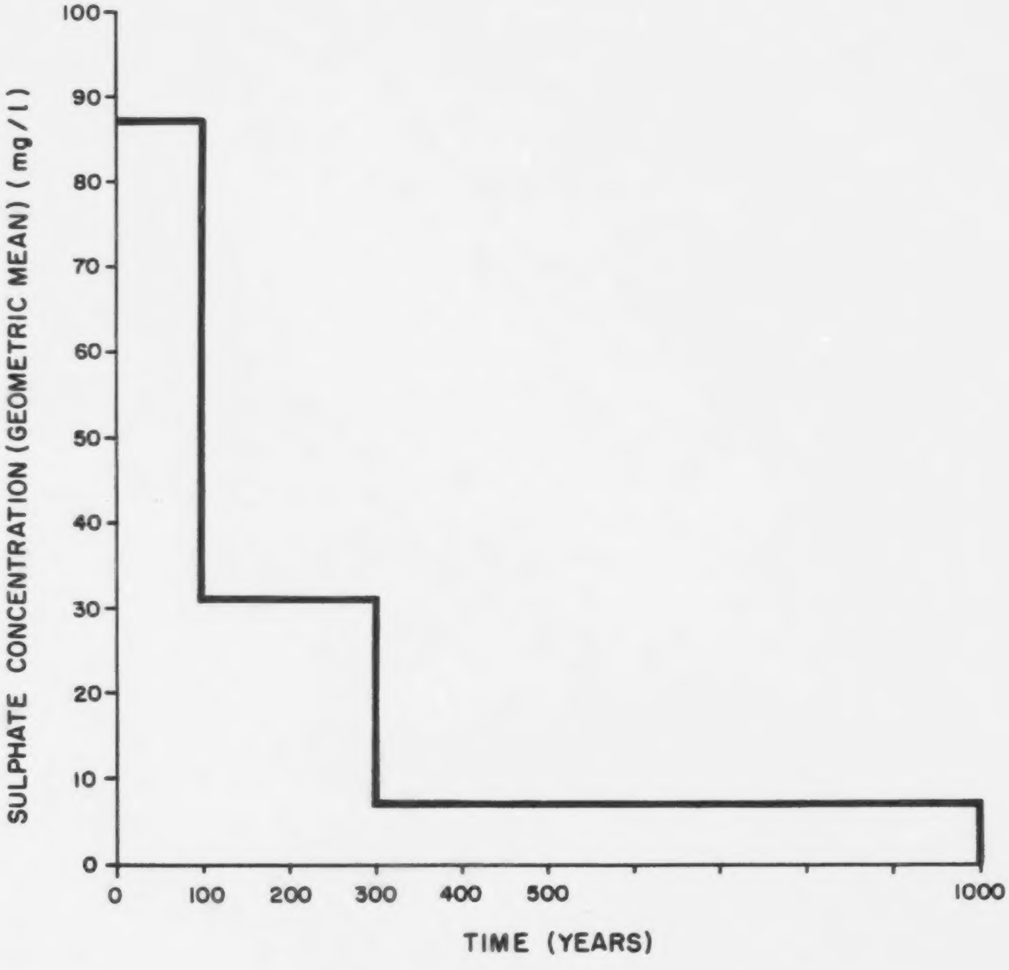

(summarized by their geometric means and geometric standard deviations) are presented in Table 4.9.

The results in Table 4.9 illustrate a situation similar to that in the previous example. The maximum sulphate concentrations resulting from a number of scenarios (3, 4, 6, 7 and 10) are in fact lower than that resulting from the null event. Since the damage of interest is the maximum sulphate concentration over the 1000 year scenario, the damage for those scenarios will, in fact, be the null event damage. The simplified event tree show in Figure 4.14 can then be constructed for this example. The percentiles of the damage distributions are presented in more detail in Table 4.10.

Again, we are at the point where the triplets for this example have been defined completely, and the uncertainty in event frequency and consequences has been quantified. Using the method outlined previously, the level 1 risk curves are constructed, combined, summarized and presented as the level 2 risk curves shown in Figure 4.15.

4.5 Example 4: A Cumulative Consequence

In this example, we consider the incorporation of a "cumulative consequence" in our risk analysis. The consequence is considered to be the 1000 year cumulative radiation dose to individuals with the characteristics of the reference receptor. It might be thought of as single year doses to 1000 persons in succession, 10 year cumulative doses to 100 persons in succession etc. A consequence such as this would likely form the basis for the calculation of a "cumulative population dose", a measure often required for optimization (ALARA) considerations.

Many of the simplifying assumptions applied in the previous examples are not appropriate when a cumulative consequence is being considered. For instance, we previously ignored all scenarios involving either multiple occurrences of the same event or the occurrence of more than one type of event. When considering a cumulative consequence, all occurrences of all events add some

Table 4.9

SUMMARY OF TIME-VARYING DAMAGES (Maximum Sulphate Concentration)

Scenario #	Description	Max. Sulphate Conc.(mg/l) GM	GSD
1	Null	87.7	1.07
2	Fire 1*	140.	1.07
3	Fire 2	49.8	2.03
4	Fire 3	11.3	1.01
5	Drought 1	161.	1.07
6	Drought 2	57.2	2.03
7	Drought 3	13.0	1.01
8	Flood 1	865.	1.07
9	Flood 2	306.	2.03
10	Flood 3	69.8	1.01

* Refers to event in time period 1,2 etc.

Table 4.10

DISTRIBUTION OF CONSEQUENCES (Maximum Sulphate concentration (mg/L))

	Scenario	10th	30th	50th	70th	90th
(1)	Null	81.5	84.2	87.8	91.8	95.5
(2)	Fire 1	128.6	135.2	140.5	145.2	153.0
(5)	Drought 1	148.0	156.0	162.0	167.0	176.0
(8)	Flood 1	794.0	835.0	865.0	895.0	940.0
(9)	Flood 2	120.0	205.0	305.0	415.0	730.0

FIGURE 4·14

SIMPLIFIED EVENT TREE: MAXIMUM SULPHATE CONCENTRATION

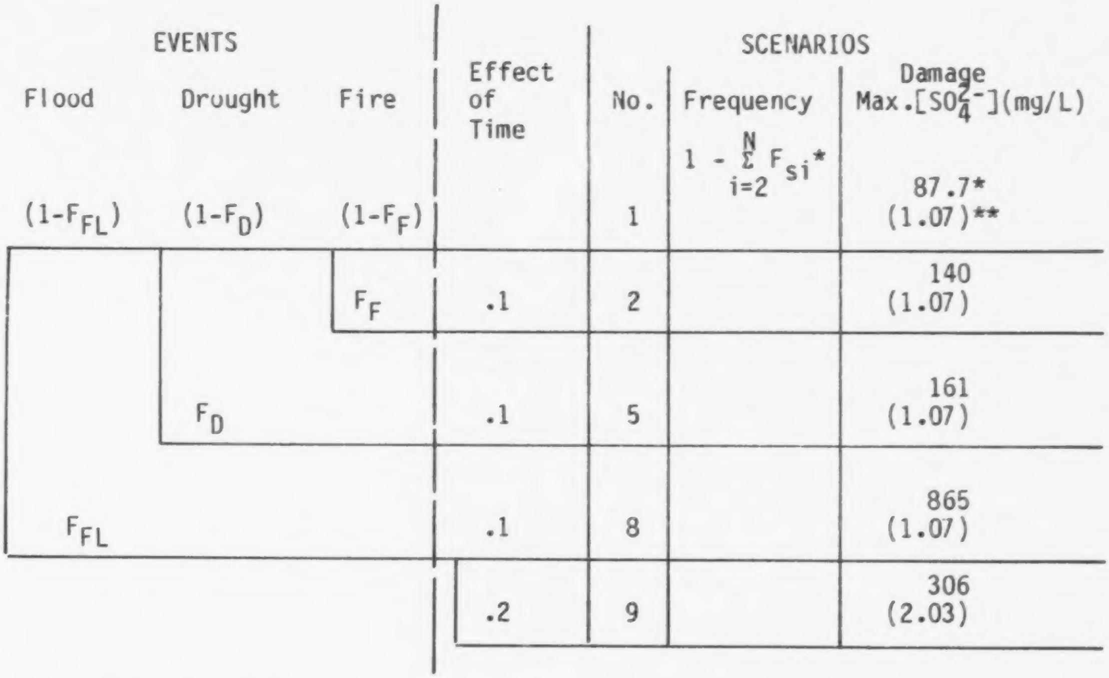

EVENTS Flood	 Drought	 Fire	Effect of Time	SCENARIOS No.	 Frequency	 Damage Max.[SO_4^{2-}](mg/L)
$(1-F_{FL})$	$(1-F_D)$	$(1-F_F)$		1	$1-\sum_{i=2}^{N} F_{si}$*	87.7* (1.07)**
		F_F	.1	2		140 (1.07)
	F_D		.1	5		161 (1.07)
F_{FL}			.1	8		865 (1.07)
			.2	9		306 (2.03)

* Geometric mean

** Geometric standard deviation

FIGURE 4·15

EXAMPLE 3 : RISK CURVE

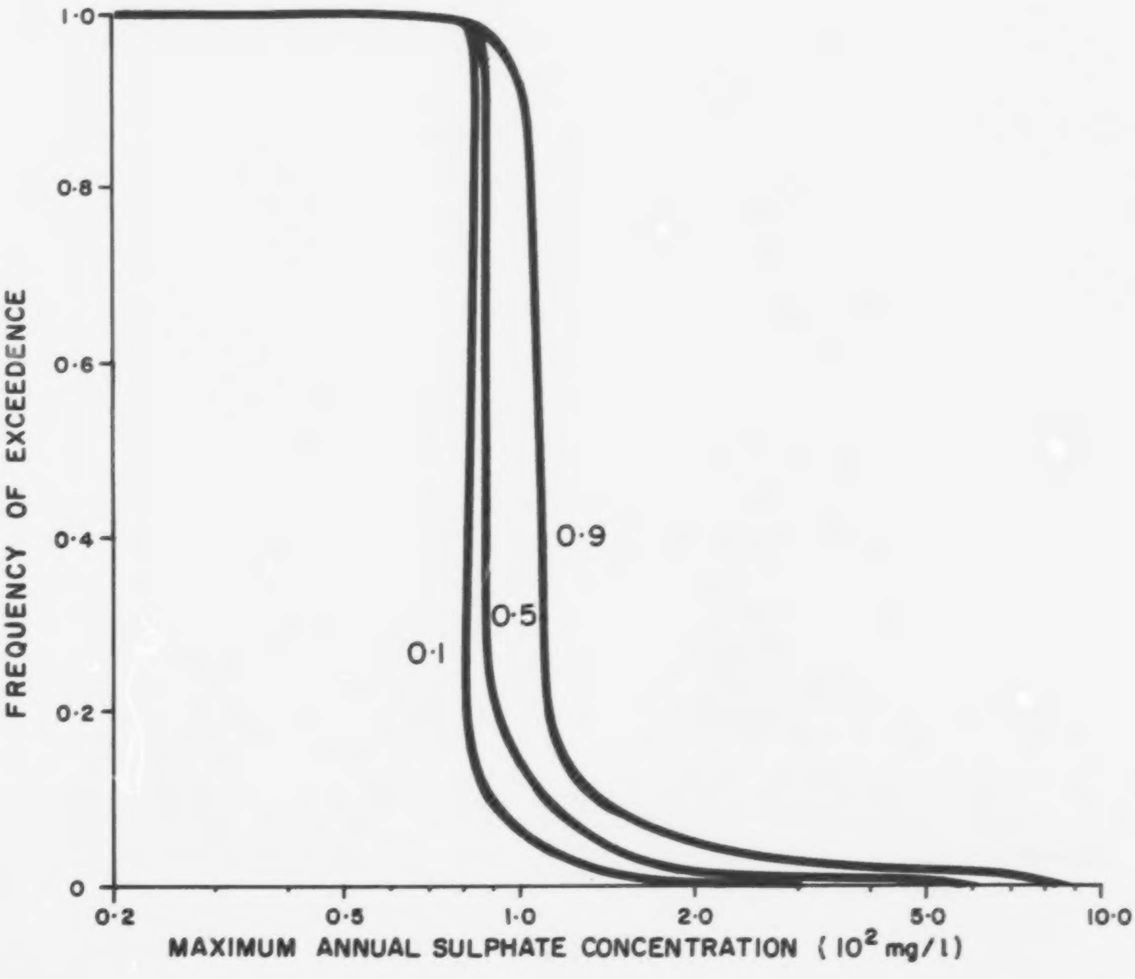

incremental damage to the calculations. Therefore the appropriate event tree potentially has many, many more scenarios than those considered in the previous examples. The fact that the null event annual dose, varies with time also continues to complicate the analysis.

For this example, the 1000 year modelling period is again divided into three windows in time, as in the previous examples. It can be shown that the incremental effects of both drought and forest fire do not vary significantly with the time of event occurrence. However the incremental effect of the flood event does depend on when the event occurs. This is because we assume that the flood event results in a dam failure and movement of some (10%) of the tailings mass into the downstream lake. The characteristics of both the solid and liquid phases of the tailings, and hence the environmental and human impact, depend on the time of occurrence.

Based on this first simplifying assumption, we might construct the event tree which is partially illustrated in Figure 4.16. As indicated on the figure, we have explicitly included scenarios involving up to 6 forest fires, and lumped those involving 7 or more into the 7+ category. Since the annual frequency of occurrence of a forest fire is relatively large compared to that for either a drought or flood event, we expect relatively more forest fires to occur in the simulation period than either drought or floods. Indeed, the expected number of droughts and floods in the simulation period is less than one, while the expected number of forest fires could be as great as seven with the assumed ranges of annual frequency of occurrence. The event tree shown in Figure 4.16 reflects this fact. The time varying incremental effect of the flood event has been included by considering scenarios involving floods in each of the three time windows. The complete event tree of this form contains 64 scenarios.

Since we felt that this event tree was somewhat unwieldy for an illustrative example, we examined further simplifying assumptions. It was shown previously (SENES, 1985b) that the incremental effects of individual forest fires was quite small, relative to the cumulative 1000 year dose. For this reason, and

FIGURE 4·16

POSSIBLE EVENT TREE FOR A CUMULATIVE CONSEQUENCE

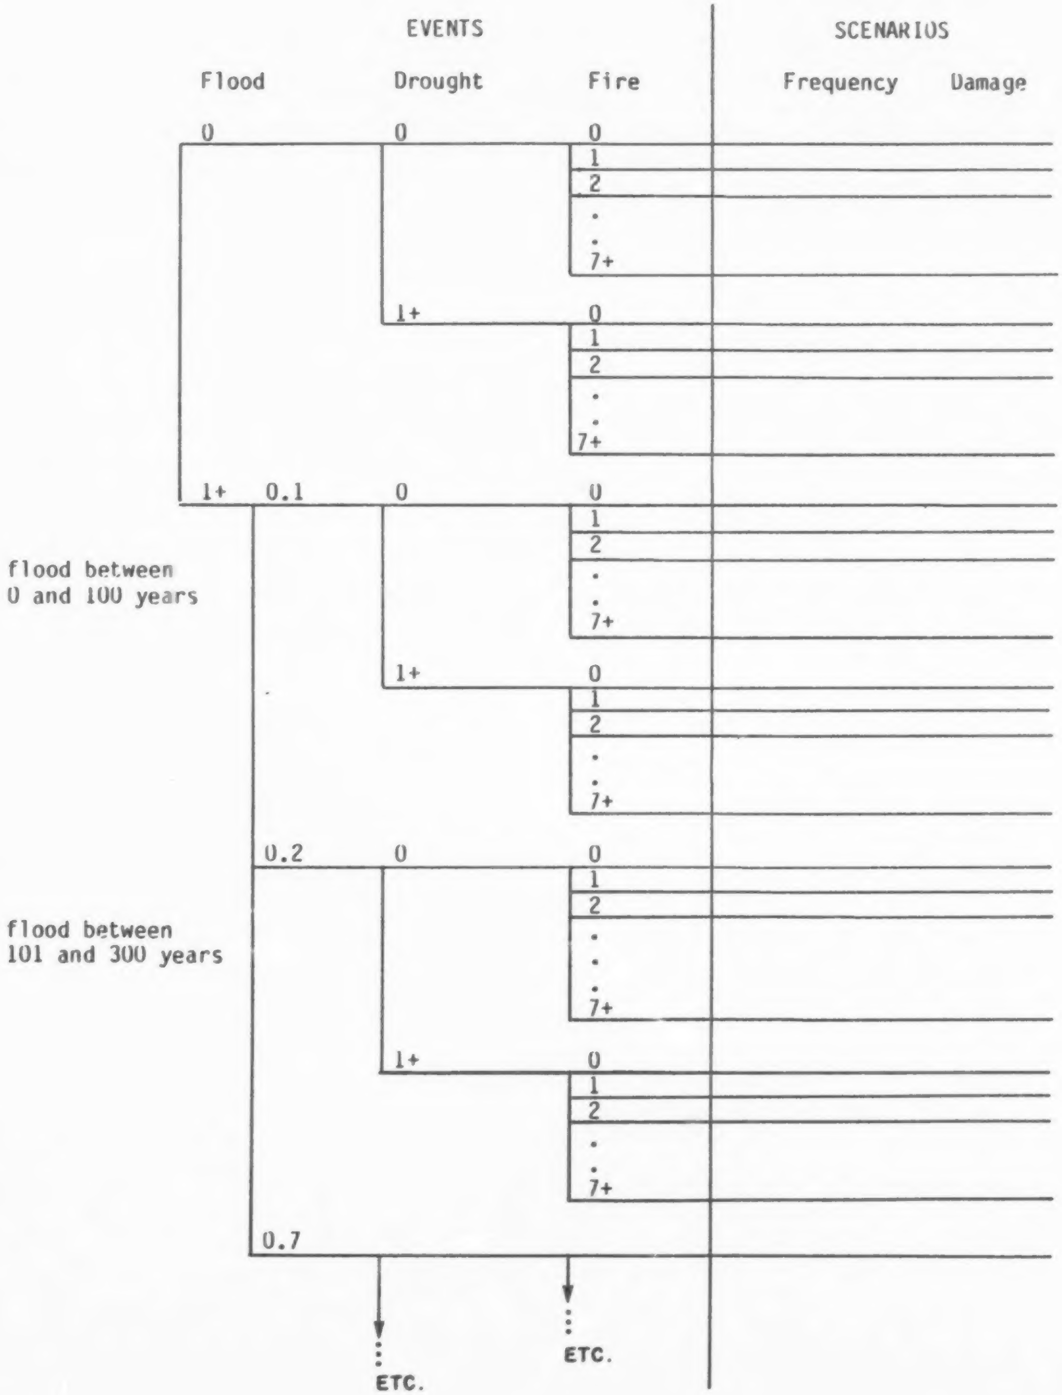
EVENTS
SCENARIOS
Flood
Drought
Fire
Frequency
Damage
0
0
0
1
2
7+
1+
0
1
2
7+
1+
0.1
0
0
1
2
7+
flood between
0 and 100 years
1+
0
1
2
7+
0.2
0
0
1
2
7+
flood between
101 and 300 years
1+
0
1
2
7+
0.7
ETC.
ETC.

since the large number of forest fires in the event tree may contribute to its complexity, it was decided to lump the multiple forest fire events into a single category. It was decided that this category would actually consist of the average number of forest fires expected over the 1000 year period. This "expected number" is easily determined from the annual frequency of event (forest fire) occurrence. The event tree resulting from this simplifying assumption is illustrated in Figure 4.17.

The event likelihoods are uncertain and hence were treated in the same manner as in the previous example. The event consequences were again based on the results of previous studies (SENES, 1985a,b). The 1000 year cumulative dose for the null event is simply the area under a curve such as that shown in Figure 2.5.

Incremental doses were determined for single fire, drought and flood events; the latter in each of the three time windows. These were then added to the cumulative null event dose, in combinations appropriate for each scenario. The number of fire events added to those scenarios involving fires depended on the annual frequency of fire, selected from the distribution of frequencies. The "level 2" risk curves constructed in this way are presented in Figure 4.18.

FIGURE 4·17

SIMPLIFIED EVENT TREE FOR EXAMPLE 4

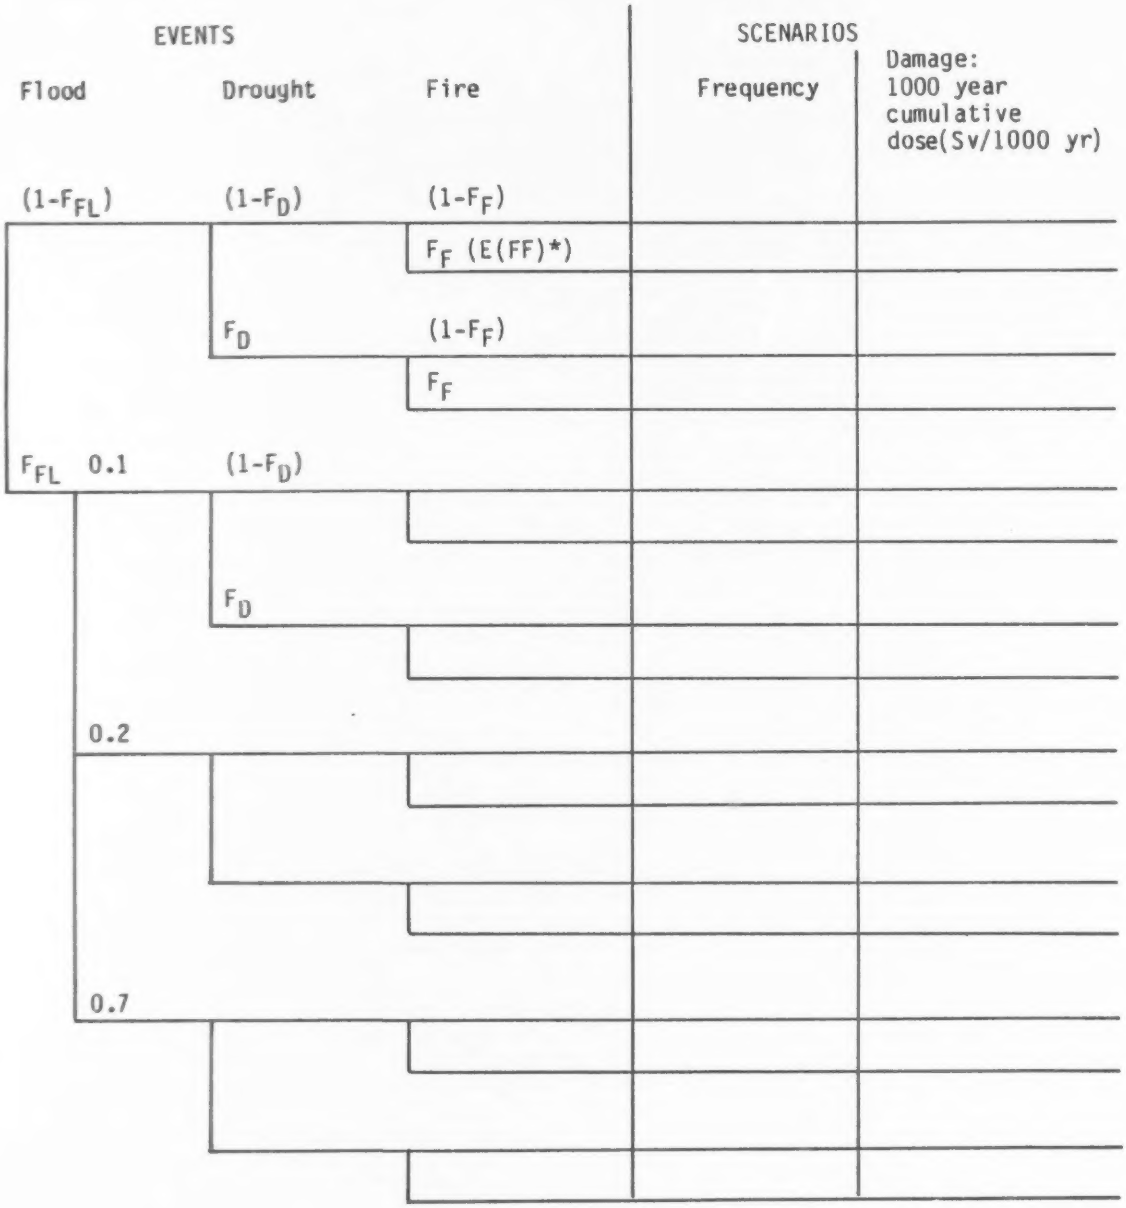
EVENTS
SCENARIOS
Flood
Drought
Fire
Frequency
Damage: 1000 year cumulative dose(Sv/1000 yr)
$(1-F_{FL})$
$(1-F_D)$
$(1-F_F)$
F_F (E(FF)*)
F_D
$(1-F_F)$
F_F
F_{FL} 0.1
$(1-F_D)$
F_D
0.2
0.7

*E(FF) The expected number of forest fires in 1000 years is used in determining consequence and is calculated as 1000 years x (annual frequency of occurrence).

FIGURE 4·18

EXAMPLE 4: RISK CURVE

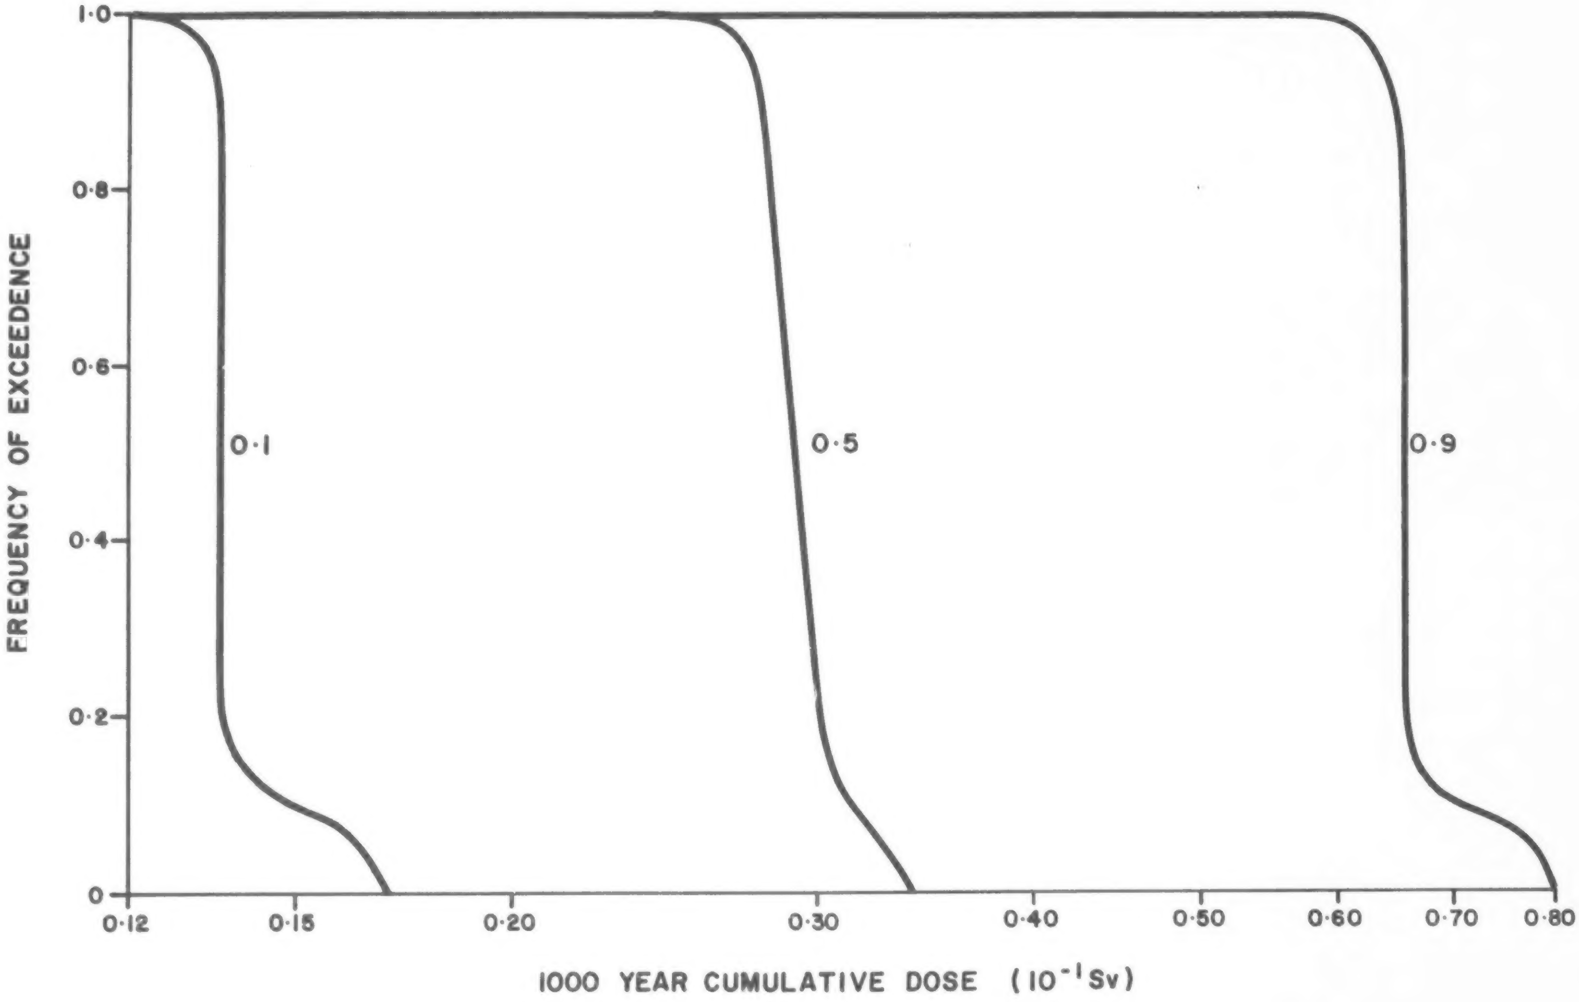

5.0 PRESENTATION AND INTERPRETATION OF RESULTS

In order to evaluate the potential risks associated with uranium mill tailings it is necessary to make predictions using mathematical models. However, as discussed previously in this report, mathematical models, by their very nature, are, at best, imperfect reflections of a real system. Moreover, the values of parameters needed to evaluate the model are often either imperfectly known or inherently variable. As a consequence, the results of all such calculations are uncertain. Since predictions into the future are required to assess the potential risks from uranium mill tailings it is not possible, even in principle, to eliminate all of the essential uncertainty by say, extensive experimentation. Rather decisions, concerning for example, the method of closing out of a tailings area, will have to be made in the face of uncertainty.

Probabilistic analysis is a tool which can be used to help quantify some of the uncertainty associated wtih the potential long-term effects of uranium mill tailings. The development of a probabilistic assessment tool has been the focus of other studies funded by the National Uranium Tailings Program (SENES 1984, 1985b). The present risk study, extends the work of a prior study (SENES 1985a) to include uncertainty about both the likelihood of a particular scenario occurring and about the consequences of that scenario, should it occur. From previous sections of this report it is clear that the techniques of probabilistic risk assessment are complex and perhaps somewhat intimidating to those not familiar with the concepts of probability or mathematical modelling. Even to persons familiar with these concepts it is not always easy to incorporate the results of a risk assessment into a decision making framework. Thus, issues related to the way in which the results of probabilistic analyses are presented, and the role of probabilistic analysis in decision making are important. Discussion of these topics form the basis for this chapter.

In large part the materials presented in Chapter 5 are based on a paper prepared by Professor M. Granger Morgan. The complete text of Dr. Morgan's paper, prepared specifically for the present study, is provided in Appendix A.

5.1 Communicating Probabilistic Results

This report adopts the concept of risk as a set of triplets:

$$R = \langle S_i, l_i, x_i \rangle \qquad (5\text{-}1)$$

where

S_i = denotes a "scenario" (i.e. what can happen)
l_i = indicates the likelihood that scenario i will occur
x_i = is a damage index, i.e. a measure of the consequences of scenario i

Since quantitative probabilistic risk assessments can involve a number of possible scenarios, and since both the likelihood and consequences of any scenario can be uncertain, we are faced with problems in understanding and communicating the results of such studies. Let us first consider some ways of presenting the results of a probabilistic analysis, in which we are certain about the scenario to be modelled.

Before examining probabilistic results, it is useful to review the fundamental features of probabilistic analysis (see Figure 5.1). The technique involves a mathematical model of some physical system. In the present situation, a uranium tailings pile and the surrounding environment are modelled. Most of the input variables required by such a model are not known with certainty. The uncertainty arises because of normal experimental error in their measurement or because (sometimes) actual measurements are impossible and judgement must be used. Also, some input variables such as wind velocities are inherently variable with time. Consequently, it is impossible for most input variables to be adequately characterized by any single number. Fortunately, as discussed previously, our full state of knowledge about any input variable can be described in the form of a subjective probability distribution.

FIGURE 5.1

CONCEPT OF PROBABILISTIC ANALYSIS

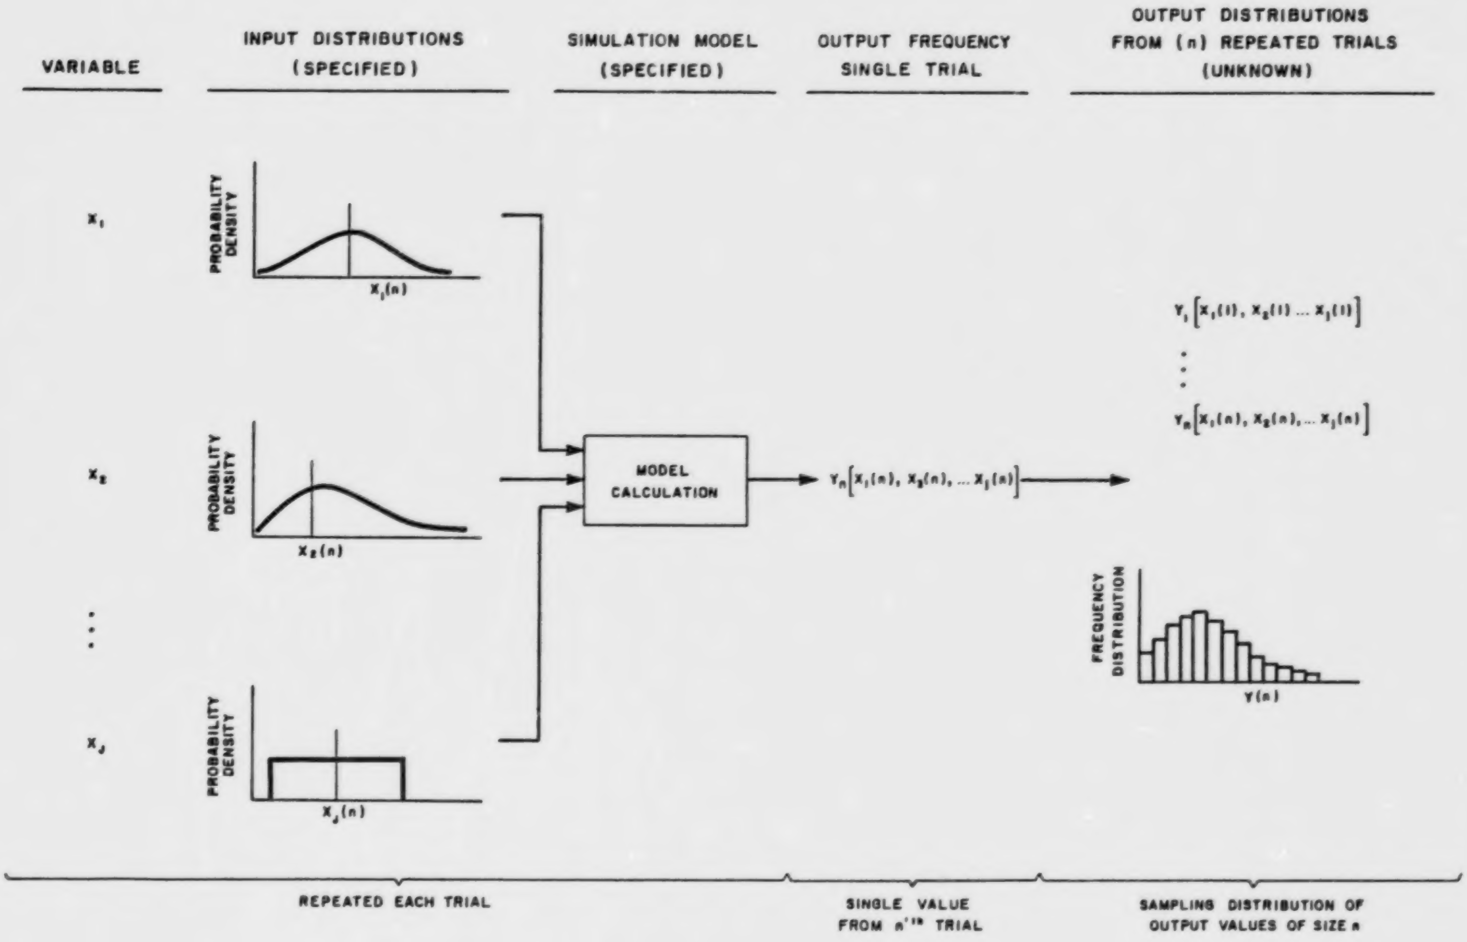
VARIABLE
INPUT DISTRIBUTIONS (SPECIFIED)
SIMULATION MODEL (SPECIFIED)
OUTPUT FREQUENCY SINGLE TRIAL
OUTPUT DISTRIBUTIONS FROM (n) REPEATED TRIALS (UNKNOWN)
X_1
X_2
. . .
X_J
PROBABILITY DENSITY
$X_1(n)$
PROBABILITY DENSITY
$X_2(n)$
PROBABILITY DENSITY
$X_J(n)$
MODEL CALCULATION
$Y_n[X_1(n), X_2(n), \ldots X_J(n)]$
$Y_1[X_1(1), X_2(1) \ldots X_J(1)]$
. . .
$Y_n[X_1(n), X_2(n), \ldots X_J(n)]$
FREQUENCY DISTRIBUTION
$Y(n)$
REPEATED EACH TRIAL
SINGLE VALUE FROM n^{TH} TRIAL
SAMPLING DISTRIBUTION OF OUTPUT VALUES OF SIZE n

Having obtained the necessary input distributions, the mathematical model is evaluated repeatedly, with new values randomly selected for each input variable on each trial. The output variable(s) from such an analysis can be considered as a sample(s) on an objective random variable(s) corresponding to that output variable(s). From the sample, the mean, variance, probability density function (pdf) or cumulative distribution function (cdf) may be estimated as desired. This information may be interpreted as describing the output variable in terms of subjective probability; telling us quantitatively what we are entitled to believe about the output variable.

The most important output from a Monte-Carlo calculation is often the probability density function. However, the cumulative distribution function is also useful in displaying certain kinds of information about the output variable. A representative pdf and cdf are shown on Figure 5.2. While both representations of the output variable contain the same information (the cdf is simply the integral of the pdf) the two representations are most useful for illustrating different characteristics of the output variables. In general terms the pdf is superior for communicating the qualitative features of the output distributions:

- the relative probability of specific values
- the most likely value (i.e. the mode) (see Appendix A, Figure 4)
- a crude estimate of the mean
- the shape of the distribution, skewness, kurtosis, etc. (see Appendix A, Figure 4)
- small changes in probability density (see Appendix A, Figure 5)

In contrast, the cdf is best for communicating the numeric characteristics of the output distribution:

- fractiles, including the median
- the probability of intervals, including confidence intervals
- tail areas
- stochastic dominance (see Figure 5.3)

FIGURE 5.2

EXAMPLE OF A PROBABILITY DENSITY FUNCTION (PDF) AND A CUMULATIVE DISTRIBUTION FUNCTION (CDF)

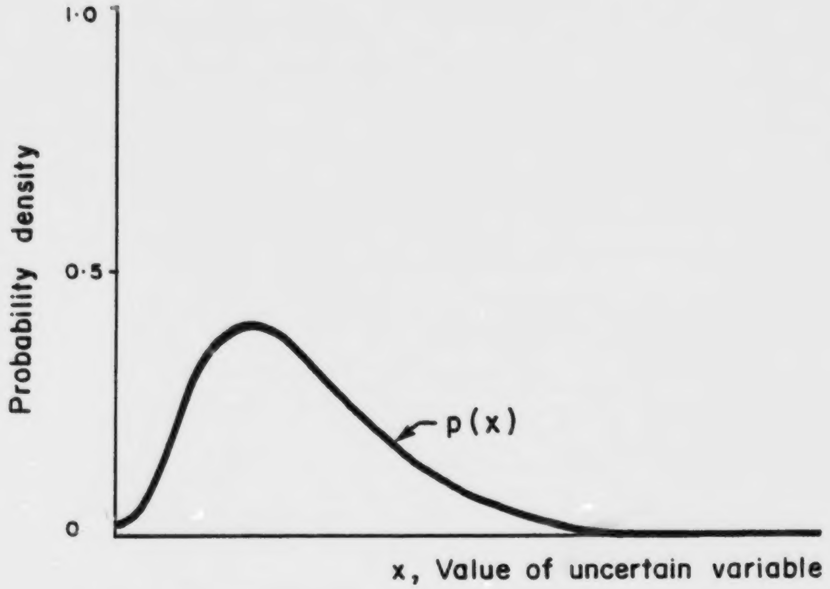

A) PROBABILITY DENSITY FUNCTION

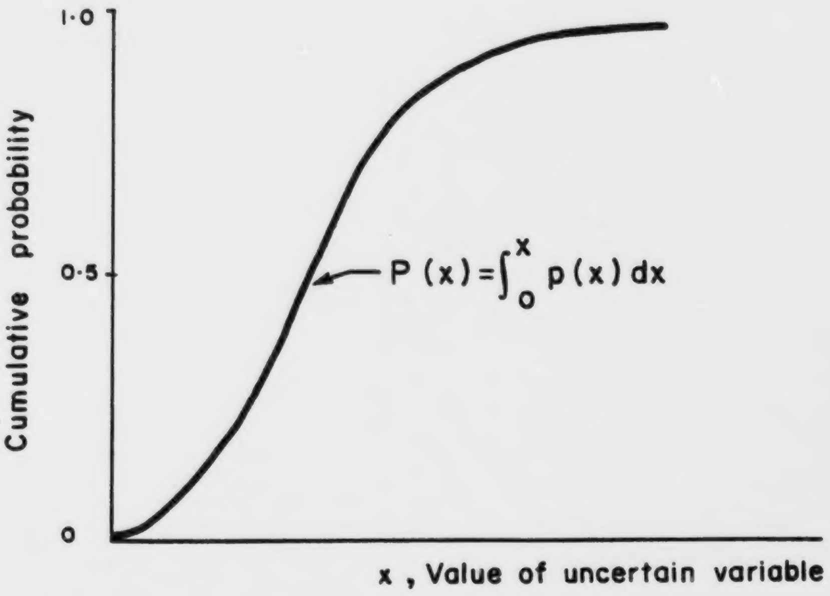

B) CUMULATIVE DISTRIBUTION FUNCTION

FIGURE 5.3

ILLUSTRATION OF STOCHASTIC DOMINANCE

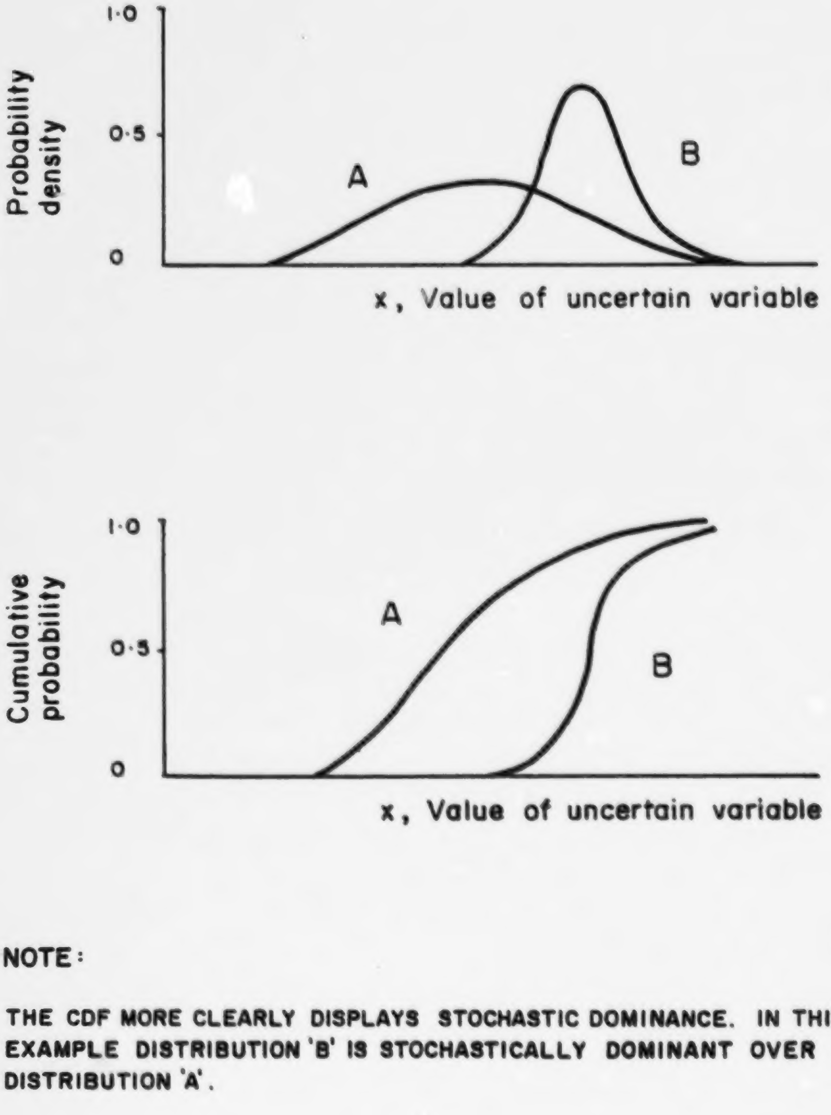

NOTE:

THE CDF MORE CLEARLY DISPLAYS STOCHASTIC DOMINANCE. IN THIS EXAMPLE DISTRIBUTION 'B' IS STOCHASTICALLY DOMINANT OVER DISTRIBUTION 'A'.

Thus, since the cdf and pdf are better at communicating different features of the same distribution, Morgan suggests it is best, when space allows, to publish both together, ideally one above the other, with the same horizontal scale as illustrated in Figure 5.2.

Most quantitative risk assessments involve a large number of uncertain quantities, as well as predicted results which vary in time and space. As a consequence Morgan (writing in Appendix A) suggests that "perhaps the most serious problem one faces in understanding and communicating the results of such studies is the problem of dealing with this high dimensionality. With one or two dimensions one can produce plots of probability density functions or cumulative distributions. When it becomes important to look at a problem across more than one or two uncertain variables, one is faced with the problem of abstracting these dimensions by "projecting" them in some way onto a two or possibly three dimensional display". Morgan also considers some examples of such projections. In this chapter we restrict our discussion to displays for two uncertain variables.

The simplest way in which to deal with two dimensions is to discretize one of them, that is to display a series of one dimensional pdf's or cdf's in the first variable conditional on the value of the second variable. An example is provided in Figure 5.4. This might represent distributions of receptor dose at several different points in time or at several different values of an important input parameter.

An alternative, of course, is to display a perspective view of the full three dimensional surface of the joint pdf or cdf with the two uncertain variables displayed on the two horizontal axes and probability density or cumulative probability displayed on the vertical axis. While this display can be helpful for providing a general sense of the results, it can be awkward to try to scale actual values from the plots. Morgan also notes that, in some cases, there may be problems in finding a viewing angle which adequately conveys the surface.

FIGURE 5.4

EXAMPLE OF TWO UNCERTAIN VARIABLES

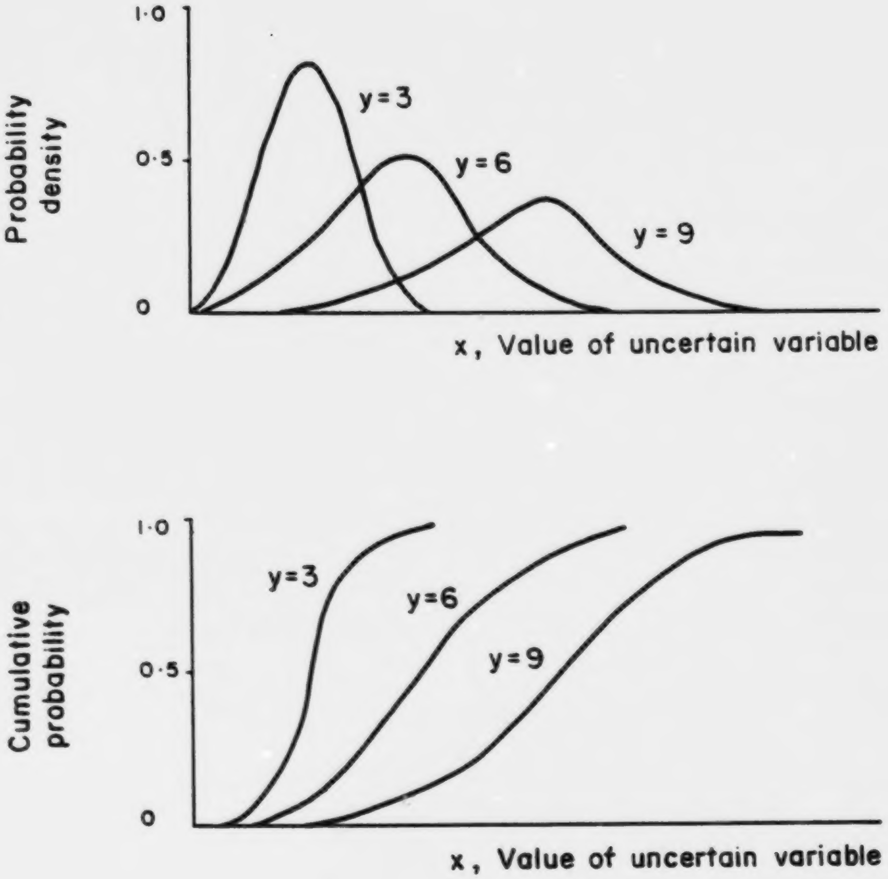

NOTE:

TREATMENT OF TWO DIMENSIONS (x AND y) BY TREATING ONE AS CONTINUOUS (x) AND DISCRETING THE SECOND (y)

In discussing the more general case, when one has a large number of uncertain quantities in a model, or even a number of alternative possible model forms, Morgan suggests that "there is clearly no way in which all of the uncertainties can be represented in a single plot".

Faced with this problem of high dimensionality one can describe two alternative extremes. The first, which we can call the "all combinations" strategy involves discretizing all of the distributions (e.g. into high, medium and low) and then running the model for all possible combinations. The typical result is mass of totally indigestible tabular displays. The opposite extreme, which we will label "naive decision analysis" treats all the variables in their full probabilistic form, adds the decision maker's values and attitudes toward risk in the form of one or more (possibly multi-attribute) utility functions and uses the standard techniques of Bayesian decision analysis to integrate across everything and produce a single "answer", which is the policy action with the highest expected net present value or utility. Morgan uses the term naive because (he claims) no good practicing decision analyst would actually adopt the approach described. In contrast to the all combinations approach which drowned us in so much information about the structure of the problem that it was unmanagable, the second approach gives us no information at all about the structure of the problem, it only tells us what to do. It provides no explanation of why this is the best choice, and no insight about how robust this conclusions may be. Clearly the objective must be to strike some appropriate balance between these two extremes. (see Appendix A).

5.2 Interpreting Risk Results

In previous sections of this report it was explained that the information about risk contained in a set of triplets can be presented in the form of an event tree (eg. Figure 1.1) or in the form of a risk curve (eg. the level 1 risk curve in Figure 1.3). Each representation contains the same information about the various scenarios, likelihoods and consequences that constitute the

FIGURE 5.5

TWO REPRESENTATIONS
OF A LEVEL 2 RISK CURVE

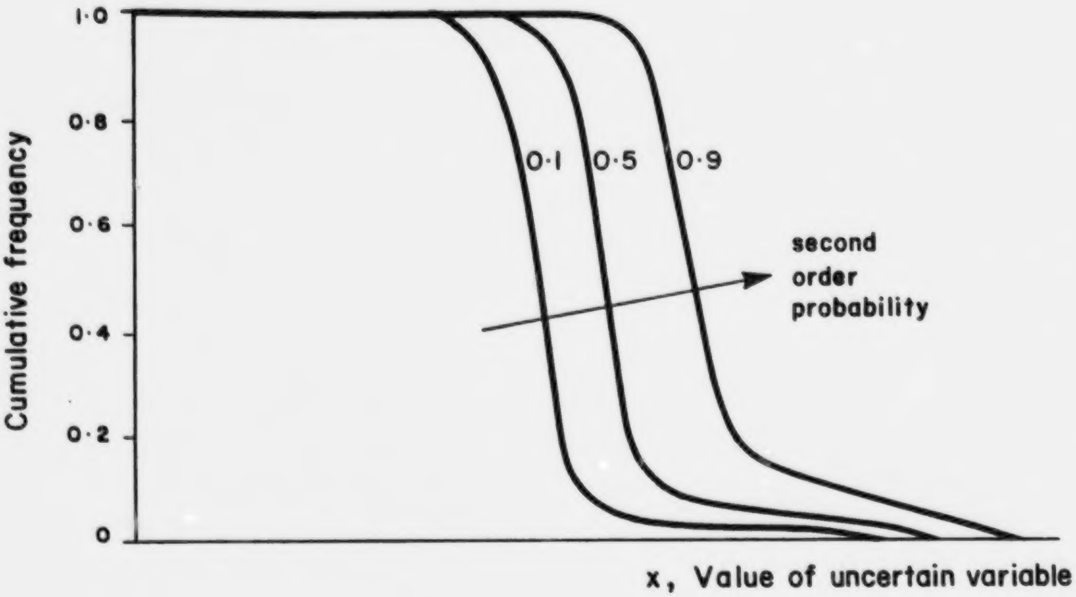

A) TYPICAL LEVEL 2 RISK CURVE

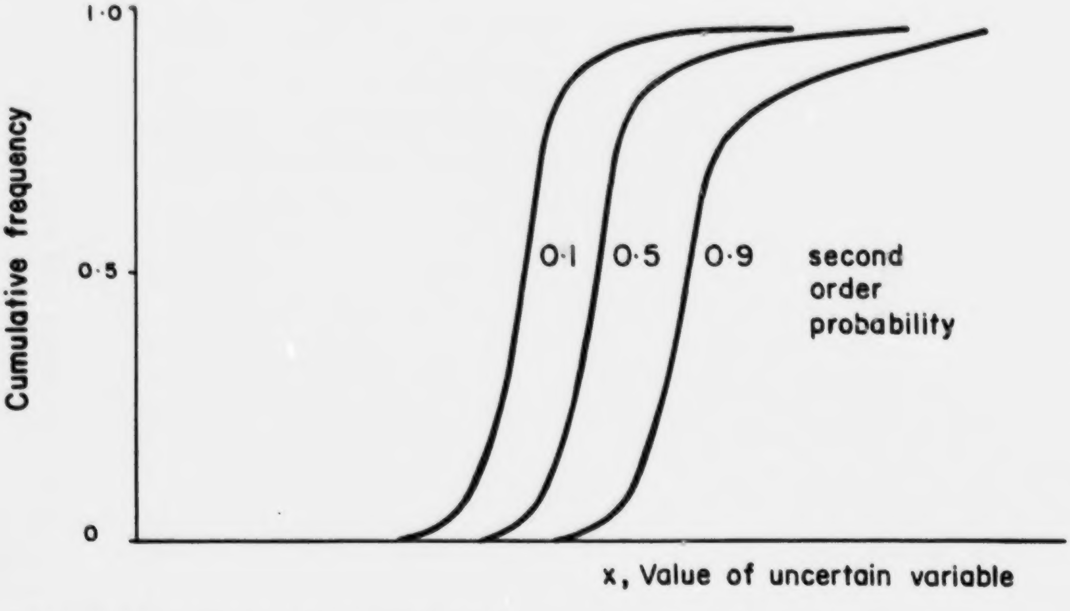

B) REPLOTTING OF FIGURE 5.5A TO MAKE IT CONSISTENT
WITH CONVENTION USED IN THIS SECTION AND APPENDIX A

risk. However, the situation is more complex than indicated by Figures 1.1 or 1.3 since:

1) the consequences of a particular scenario are uncertain and can be represented as pdf's or cdf's as described in the previous section and
2) the frequency (i.e the likelihood that a particular scenario will occur) is also uncertain.

Following the terminology of Kaplan and Garrick (1981) we express our uncertainty about the frequency in a "probability of frequency" format as illustrated, for example, in the level 2 risk curves shown in Chapter 4. Morgan (writing in Appendix A) refers to this as uncertainty about uncertainty or alternatively as second level uncertainty.

In earlier sections of this report the level 2 risk curves are shown as downward or complementary cumulative frequency curves such as that shown in Figure 5.5(a). In Figure 5.5(b) the level 2 risk curve is replotted to be consistent with earlier discusssions in this chapter and the conventions used by Morgan in Appendix A (i.e. complementary cdf changed to cdf).

Following the discussion in Appendix A, let us begin our discussion by considering how such a family of curves might be interpreted if they were to be presented as an expert's subjective judgement for use as the input in a risk analysis (Recall that, to this point, such curves have been the output of an analysis).

Suppose that such an expert, when interviewed using appropriate procedures for the elicitation of expert subjective judgement, produced a result such as that shown in Figure 5.6(a). Now suppose that we went on to ask this expert questions about how much the curve he has given us might change and still (more or less) represent his opinion. For example we might show him many curves, such as the spaghetti of curves illustrated in Figure 5.6(b) and ask him the extent to which he can distinguish between these curves and the curve

FIGURE 5.6

ILLUSTRATION OF CONSTRUCTION OF SECOND ORDER UNCERTAINTY IN A CDF ELICITED FROM AN EXPERT

(SEE TEXT FOR DISCUSSON)

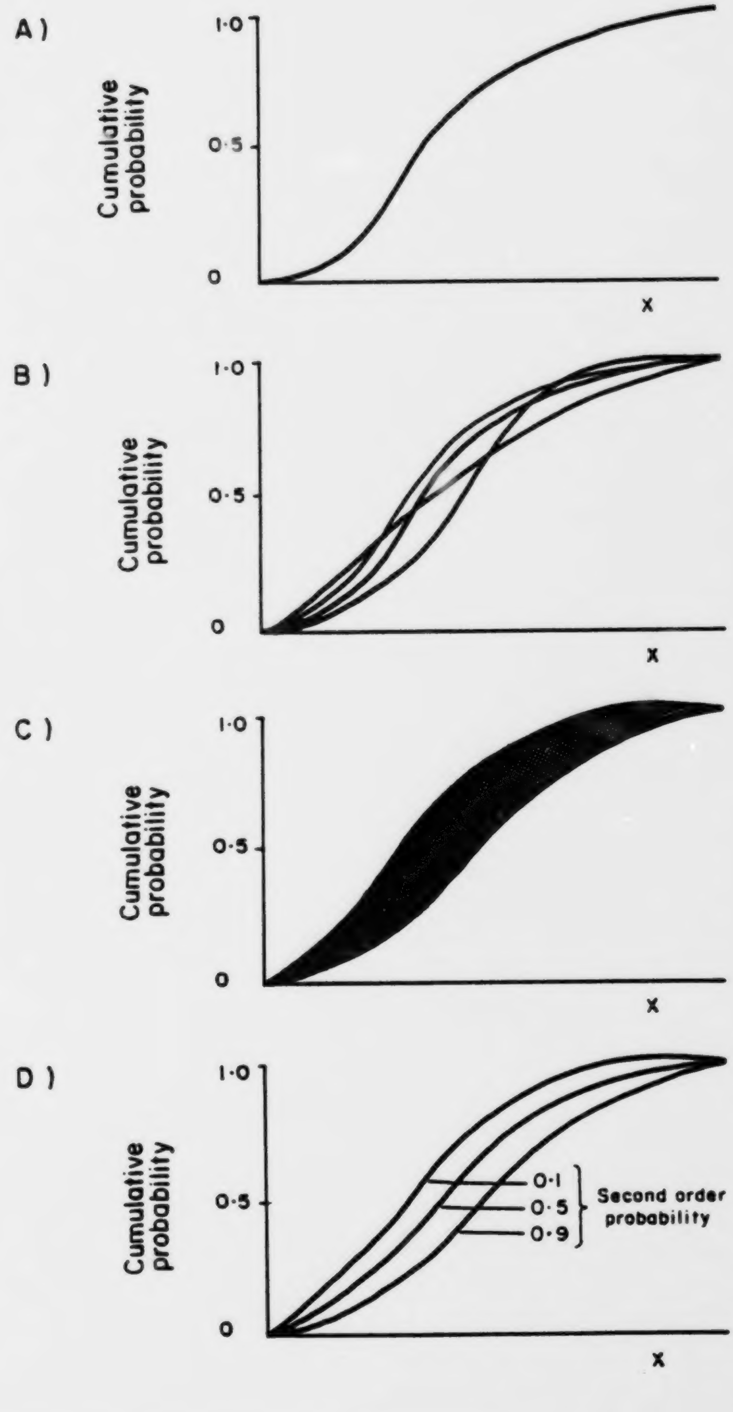

he initially provided. In most cases, the expert is likely to recognize that his judgement is probably not reliable enough to distinguish meaningfully between members of a family of such similar curves. Thus (given a variety of assumptions and some analytical manipulation which we gloss over) one might be able to represent his judgement by an envelope of curves such as that shown in Figure 5.6(c) or with further effort by a result such as Figure 5.6(d) which is a result in the same form as Figure 5.5(b) (Morgan cautions that in practice several issues could complicate such an analysis - see Appendix A).

Morgan suggests there is another way in which one could imagine constructing curves such as those shown in Figure 5.6. Suppose that one has a large number of alternative models of physical reality, each of which produces a slightly different outcome. The result might be a family of curves quite analogous to that shown in Figure 5.6(b) With some appropriate additional judgements, one could then proceed from this to produce a curve like Figure 5.6(d).

Whether it is useful to preserve information about second order uncertainty depends entirely on the intended use of the results. If, for example, the decision maker has to meet an established probabilistic limit, a decision could be made on the basis of the size of the tail distribution which exceeds the specified limit. Alternatively, a more complete probabilistic limit could be formulated, as for example by the specification of a cdf which must stochastically dominate the actual distribution generated by the analysis. That is, the calculated cdf must lie completely above and to the left of some limiting cdf. This concept is illustrated in Figure 5.7.

If, on the other hand, the decision maker has an absolute limit (a limit which must not be exceeded under any circumstances), the probabilistic format may not be as useful. This is because there will always be a low, but nevertheless finite, probability of having a consequence greater than that observed in any particular analysis (for example we would expect to find a few additional high tail values with each additional set of Monte Carlo trials). In this case it may be appropriate to integrate out the uncertain variables and present the results in a simpler form (for example the expectation of

FIGURE 5.7

THE APPLICATION
OF STOCHASTIC DOMINANCE CRITERIA
TO RISK ASSESSMENT

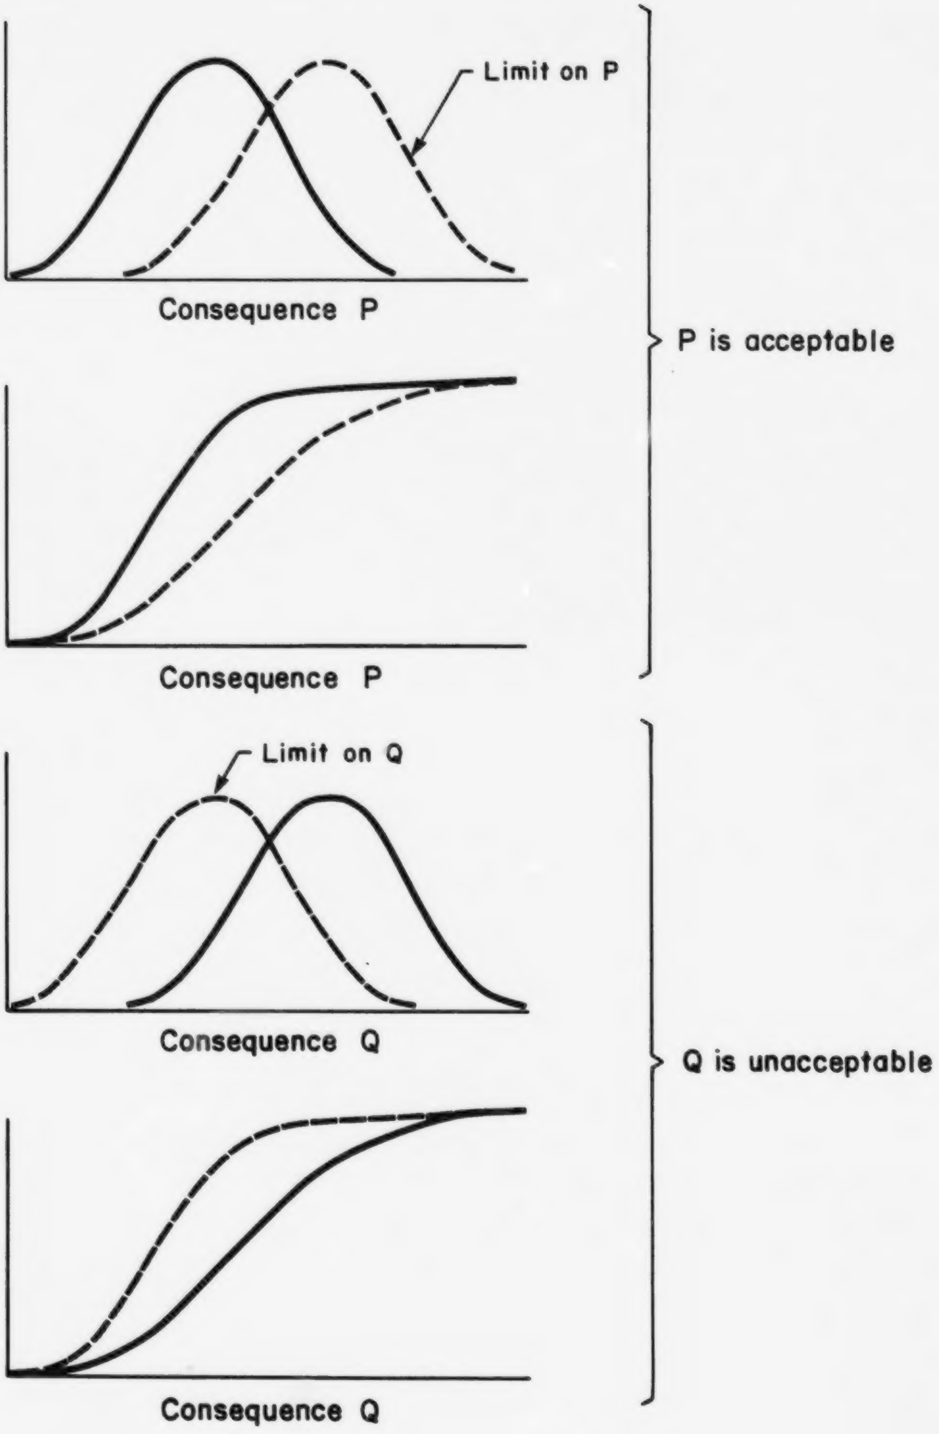

risk).

The choice depends on how the decision maker intends to make his decision and how useful the results are in providing insight to the problem.

5.3 The Role of Quantitative Risk Analysis in Risk Management Decision Making

Most people have a mental model of risk management decision making that looks something like Figure 5.8. A policy maker who is responsible for making risk management decisions identifies the problem that he or she faces and requests a group of analysts to undertake an anlysis. They set up or "frame" the problem, collect needed data and build models, do the analysis, perform sensitivity studies and produce a set of findings and "answers" which are communicated back to the decision maker. The decision maker combines these results with various other relevant considerations and reaches some decision.

While the model represented by Figure 5.8 probably accurately describes most people's view of the role of analysis, it serously misunderstands the true value of quantitative analysis in difficult risk and other decision making problems. In such problems, analysis is often better at providing insight into the nature of the problem than it is at producing "the answer". Analysis can help decision makers to refine and improve their framing of the problem and help them to better identify and justify the set of policy options they should be considering.

Figure 5.9 illustrates an alternative approach to the role of quantitative analysis in risk management decision making. As in the previous model, the process begins when the policy maker identifies a problem and turns to the analysts for assistance. At this point the activities again move from "The policy maker's world" to "the policy analysts' world".

While the policy makers will have some general thoughts about the nature of the problem, it is the analysts who will develop an initial framing of the problem, collect some initial data and build some simple models in order to do

FIGURE 5.8

THE COMMON VIEW OF THE ROLE OF ANALYSIS IN THE RISK MANAGEMENT POLICY PROCESS

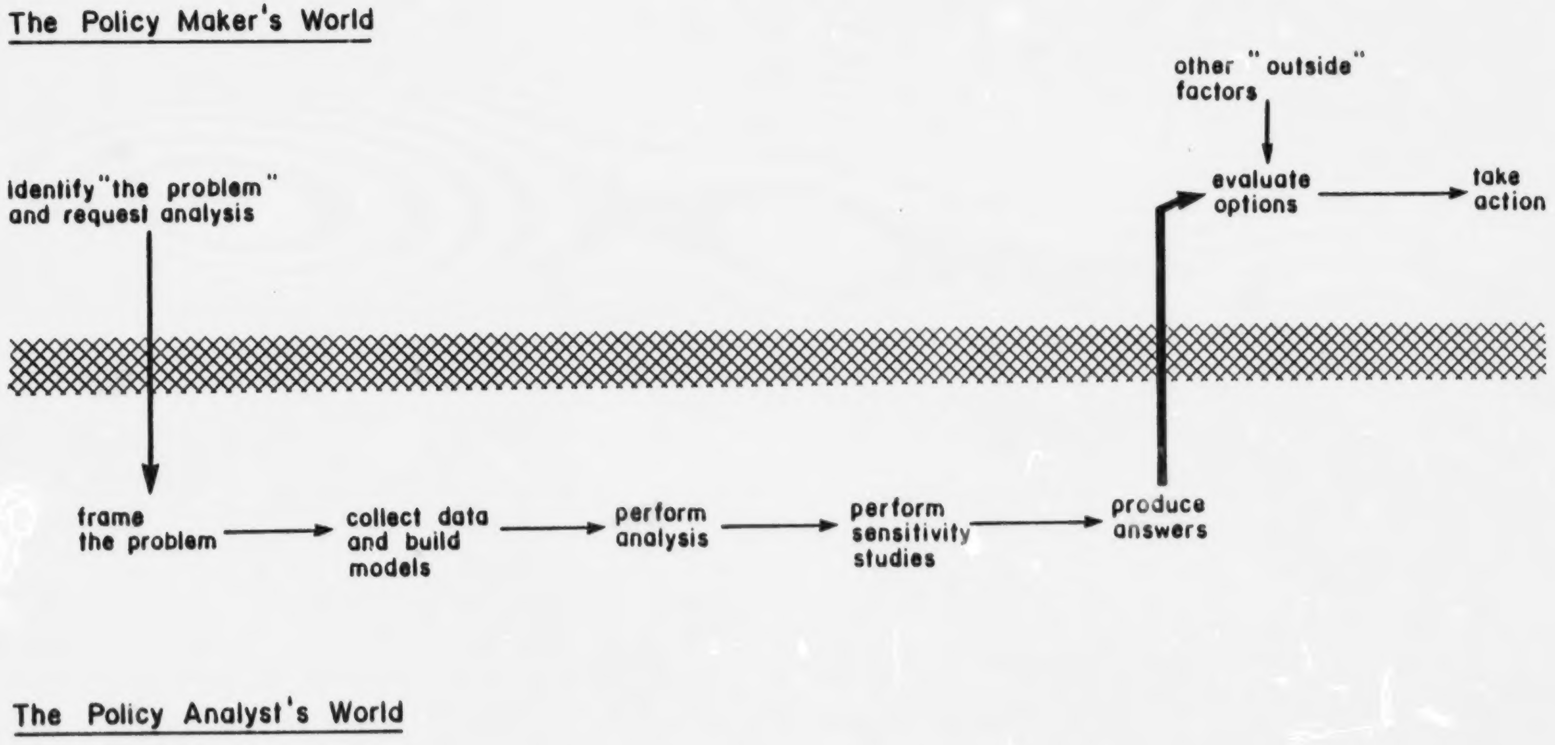
The Policy Maker's World
other "outside" factors
identify "the problem" and request analysis
evaluate options
take action
frame the problem
collect data and build models
perform analysis
perform sensitivity studies
produce answers
The Policy Analyst's World

FIGURE 5.9

AN ALTERNATIVE APPROACH
TO THE ROLE OF ANALYSIS
IN THE RISK MANAGEMENT POLICY PROCESS

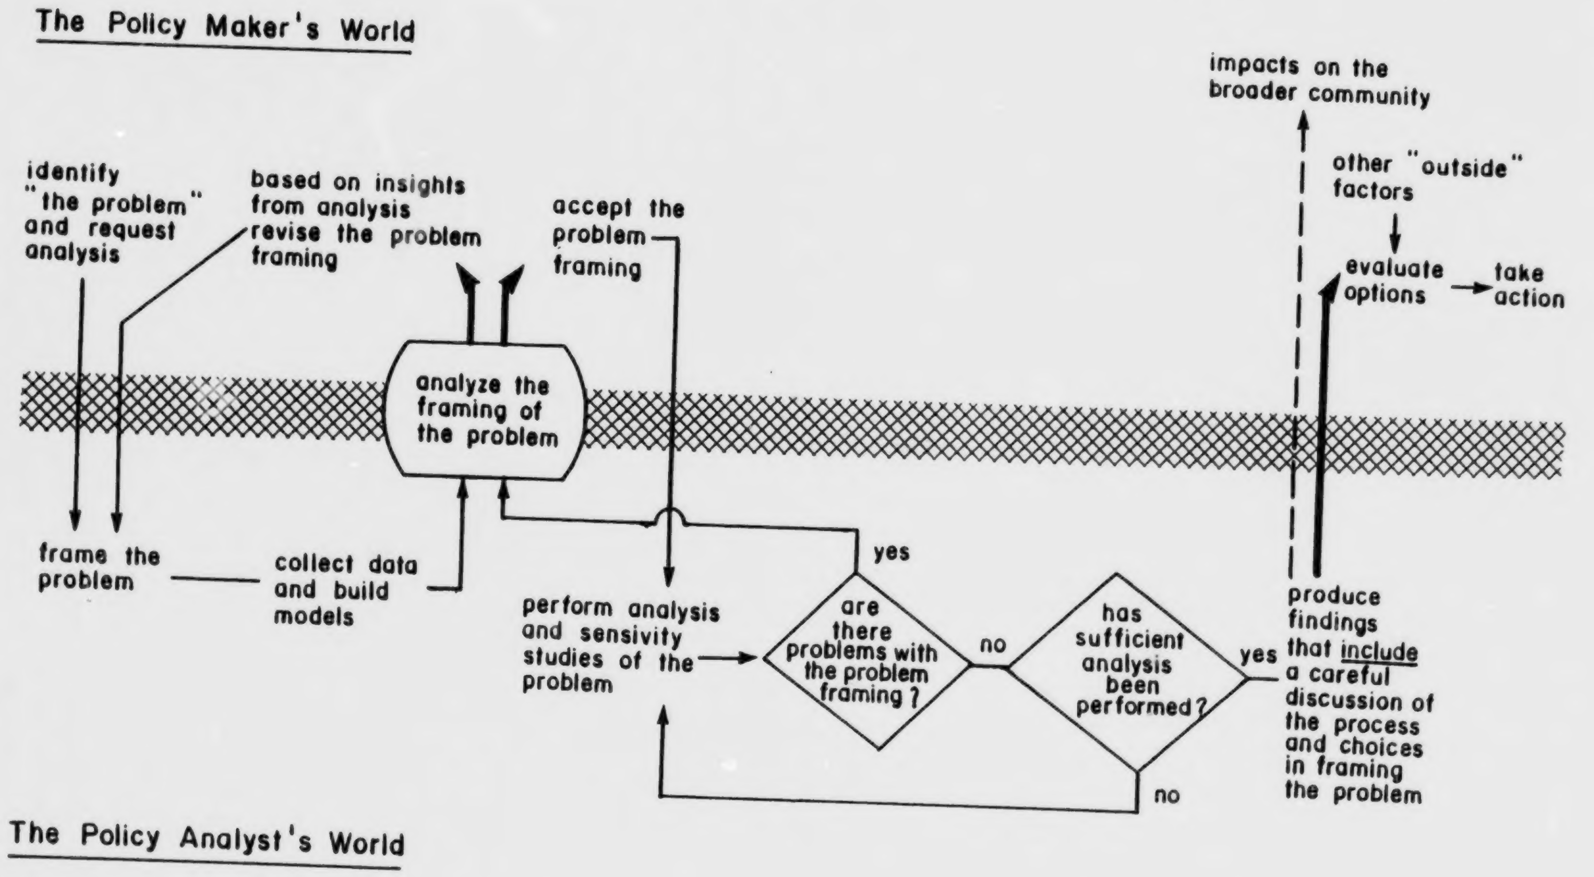
The Policy Maker's World
identify "the problem" and request analysis
based on insights from analysis revise the problem framing
accept the problem framing
impacts on the broader community
other "outside" factors
evaluate options
take action
analyze the framing of the problem
frame the problem
collect data and build models
perform analysis and sensivity studies of the problem
yes
are there problems with the problem framing ?
no
has sufficient analysis been performed ?
yes
no
produce findings that include a careful discussion of the process and choices in framing the problem
The Policy Analyst's World

preliminary analysis ... with the objective of evaluating the problem framing. That analysis allows the modellers and the decision makers to understand what the important characteristics and constraints of the problem are, what questions it makes sense to ask, what questions are essentially unanswerable, etc. After one or more iterations of this type, a revised statement of the problem is jointly developed. It then makes sense for the analysts to go on and do more detailed analysis and sensitivity studies. Sometimes in the process of these more detailed studies, problems may be uncovered which suggest that one should take another pass at reformulating the overall framing of the problem. Finally, after enough iterations, one will reach the point where a satisfactory analysis of the final problem has been produced. This is then communicated back to the decision maker, along with a description of the considerations which went into the selection of the final framing of the problem.

According to this model (described in more detail by Morgan in Appendix A) the process of framing the problem is an essential element of the risk analysis process. In particular, this iterative approach to problem framing and model building provides greater opportunity to adequately resolve a number of important issues. These include the need to:

1) develop an explicit statement of objectives for the modelling exercise that answers the question -" exactly what is it I am supposed to model?" In the case of uranium tailings, is it, for example, the concentration of radium in lake water?, the annual dose to a person living in proximity to the tailings? if dose, what pathways should I consider?: how far into the future must I model? etc.
2) define and evaluate the assumptions needed to construct the model(s)
3) establish system components and their relationships
4) identify data and evaluate data needs
5) communicate with other analysts from various disciplines as well as with the policy makers
6) develop and test hypotheses as the model development proceeds.

It is also obvious that tasks of formulating and communicating the results of various stages of the risk analysis process to the decision maker in a useful form are crucial to the risk analysis process.

Put another way, many of the benefits of risk analysis come from the problem framing and model building process itself rather than the actual results of the analysis (i.e. the specific numbers). In the context of probabilistic risk analysis this observation is very important. The kind of analysis described in this report and elsewhere (eg. SENES 1984, 1985a,b) forces us to recognize the limitations of what we know today, and the limitation and uncertainties associated wtih modelling in a way that is not often possible with more conventional modelling approacnes.

6.0 SUMMARY AND CONCLUSIONS

This study examined a number of issues related to the long-term risks associated with uranium mill tailings. The primary focus was on the development and demonstration of a methodology for incorporating time dependence and uncertainty into estimates of risk. The application of the de minimis concept to probabilistic risk assessment was explored, as were the interpretation and presentation of risk analysis results.

Many of the relevant findings are summarized in the following sections.

6.1 Application of the De Minimis Concept

The de minimis concept is applied implicitly in many individual and societal decisions. The essence of the concept is that there are risks which are so small as to be considered insignificant, or trivial. The primary advantage derived from an explicit application of the de minimis concept is that it may help to prevent the spending of society's resources (which are finite) on insignificant matters, to the detriment of those more deserving.

It has been proposed that a de minimis value might be useful in regulating the exposures resulting from uranium mill tailings. Such a value might serve as a level below which individual exposures need not be considered further. It may also be useful in identifying limits for the sizes of populations (indirectly via distance considerations) or durations of time over which collective doses would need to be calculated.

The application of the concept becomes much more difficult when a probabilistic approach is taken to estimating individual and collective consequences. Different methods, each with their own advantages and disadvantages were identified for applying a de minimis criteria to probabilistic estimates of both individual and collective dose. One approach, in examining individual consequences, is to simply ignore any probabilistic trials which produce estimates of dose below the de minimis value. A preferred approach is to examine some attribute of the distribution of annual dose to the receptor. For example, if the mean value of the distribution is less than the de minimis value, the dose might be considered trivial. One might also wish to specify some degree of confidence (say 95%) that the mean value is less than de minimis.

In applying a de minimis value to the truncation of collective dose in time and space, one could simply continue, in each probabilistic trial to calculate and sum the receptor dose until it fell below some de minimis level. However, there is the potential for the time and/or distance required, on a very few trials, to be quite large. A second approach is to calculate the distributions of dose at many points in time and space. Those distributions whose mean values fall below the de minimis value would be excluded from the calculation of collective dose.

6.2 Incorporation of Time Dependence

The physical and chemical characteristics of a uranium tailings area inevitably change with time. Thus the impacts on the surrounding environment and human population will also change with time, regardless of the occurrence of episodic events. The incremental impacts of episodic events also have the potential to vary, depending on when they occur. In any case, the radiological impact on nearby humans, in any given year, or cumulated over many years, depends on when the episodic events occur.

To incorporate this phenomena into estimates of risk, the entire 1000 year modelling period was divided into several "windows in time". The number and location of the time windows selected for any application depend primarily on how the null event dose varies with time. Time should be divided into intervals such that the null event dose does not vary greatly within each interval. That is, the dose calculated for any year within the interval should be able to be considered somewhat representative of the dose in any other year in the same interval.

6.3 Incorporation of Uncertainty

The methodology developed and demonstrated in this study incorporates the uncertainty associated with both the frequency of event occurrence and with the scenario consequences. The method depends first on the ability to construct and specify an appropriate event tree; outlining the events, scenarios and consequences to be included. The (uncertain) frequency of event occurrence must be specified as a distribution of values, as must the (uncertain) consequence of each scenario.

The uncertain event frequencies and uncertain scenario consequences are then combined to produce a number of "level 1" risk curves, each with an associated probability. These curves are then summarized; essentially by plotting equipercentile points for a number of different consequence values. The results of this operation are expressed as "level 2" risk curves; representing the frequency with which any given consequence or greater occurs, and our confidence in that estimate of frequency.

6.4 Presentation and Interpretation of Results

Probabilistic results, in general, are best summarized by presenting both the probabability density function (pdf) and cumulative distribution function (cdf) for the desired variable(s). The pdf best summarizes qualitative aspects of the distribution, while the cdf provides a better indication of its quantitative characteristics. In dealing with multi-dimensional problems, one may wish to construct three-dimensional figures, or projections of several dimensions on a two or three dimensional surface.

The results generated by the risk analysis methodology described in this report may be summarized and displayed either as event trees or risk curves (level 1 and/or level 2). The level 2 risk curves, express a "probability of frequency", or "uncertainty about uncertainty" regarding a certain consequence. These results can be interpreted quantitatively as they are presented. More intuitively, one may wish to visualize these results as the subjective judgement of an expert or group of experts, regarding the quantity of interest.

Such risk curves might be used in regulatory decision-making in several ways. For instance, a regulation might be constructed in terms of probabilistic limits; either on the size of the tail distribution or on an entire cdf. In this case, the risk curves would be directly relevant. However, if regulations are framed in terms of absolute, deterministic limits which may not be exceeded, probabilistic results may not be particularly useful. In this case, it might be appropriate to present the results in a simpler, integrated form.

6.5 Conclusions/Recommendations

The concept of risk as a set of triplets assumes that we are able to define a finite set of events and scenarios and express their likelihoods and consequences. When temporal aspects are not considered, and when uncertainty in many of the quantities is ignored, the analysis is relatively straightforward. However, the incorporation of either time dependency or uncertainty complicates the analysis considerably. When both are considered, the situation becomes quite complex.

The implications of adding time variance are that the event trees can become unmanageably large (particularly for cumulative consequences). Generating distributions of consequence for all the scenarios involves developing and implementing some very complex models, and evaluating them many times. This has the potential to become a very onerous task, potentially requiring hours of computing time. Furthermore, the incorporation of uncertainty in event frequencies has the potential to greatly complicate the calculation of scenario frequencies. In general, the methodology involves the generation and manipulation of large quantities of information. Many simplifying assumptions are required to keep the exercise tractable.

It is important that the potential advocates and/or users of the techniques demonstrated in this study realize the level of effort required to perform the analyses properly. One must then weigh this effort (or cost) against the potential benefits or insights to be gained from the analysis.

A risk analysis, such as that outlined in this report, may be most useful in providing insight into a problem, rather than in generating "the answer". As Morgan notes (in Appendix A), the point of doing analysis on risk management problems "should be to try to frame the problems we face in useful and productive ways and develop insights that will help us do a better job in risk management. The literatures of decision and policy analysis have developed a large number of analytical and display tools that can be used in this process. They are not however, ends in themselves. They should be used only if they help to give insight and understanding. The choice or which to use should depend on how helpful...after suitable practice and experience...analysts and decision makers find them."

Finally, it should be noted that the methodology described should not be applied in a rote, mechanical fashion. Analysts and regulators alike should keep in mind that risk assessment is at least as much an art as a science. Quantitative risk assessment can provide understanding and insight, clarity and communication. It is a valuable aid to good human judgement but should never become a substitute for it.

7.0 REFERENCES

Atomic Energy Control Board, 1985. "Consultative Document C-85, The Basis for Exempting The Disposal of Certain Radioactive Materials from Licensing". 6 May.

Box, G.E.P., Tiao, G.C., 1973. "Bayesian Inference in Statistical Analysis." Published by Addison-Wesley.

Chaterjee, R.M., Coady, J.R. and Wagstaff, K.P., 1982. "An Approach to the Exemption of Materials from Regulations as Radioactive Wastes." AECB paper INFO-0086, 15 September.

Kaplan, S. and Garrick, B.J., 1981. "On the Quantitative Definition of Risk." Risk Analysis, Vol. 1, No. 1, pp. 11-27.

Kaplan, S., 1981. "On the Method of Discrete Probability Distributions in Risk and Reliability Calculations-Application to Seismic Risk Assessment." Risk Analysis, Vol. 1, No. 3, pp. 189-196.

McInnis, P., 1985. "Long-term Effects of Natural Events on Uranium Tailings Impoundments." Research Report prepared for the National Uranium Tailings Program, CANMET, EMR. NUTP 85-3 (TR).

SENES Consultants Limited, 1984. "A Study to Develop and Demonstrate Techniques for the Probabilistic Assessment of the Long-Term Effects of Uranium Mill Tailings in Canada - Phase I." Research Report prepared for the National Uranium Tailings Program, CANMET, EMR, under Supply and Services Canada Contract No. 0SQ83-00236.

SENES Consultants Limited, 1985a. "Approaches to Risk Assessment for Canadian Uranium Mill Tailings." Research Report prepared for the National Uranium Tailings Program, CANMET, EMR, under Supply and Services Canada Contract No. 0SQ84-00335.

SENES Consultants Limited, 1985b. "A Study to Develop and Demonstrate Techniques for the Probabilistic Assessment of the Long-Term Effects of Uranium Mill Tailings in Canada - Phase II." Research Report prepared for the National Uranium Tailings Program, CANMET, EMR, under Supply and Services Canada Contract No. 0SQ84-00207.

SENES Consultants Limited, 1985c. "A Discussion Paper on *De Minimis* (or Trivial) Levels of Contaminant Concentrations or Doses." Research Report prepared for the National Uranium Tailings Program, CANMET, EMR, under Supply and Services Canada Contract No. 26SQ84-00198.

United States Nuclear Regulatory Commission (USNRC), 1986. "10 CFR Parts 19, *et al*, Standards for Protection Against Radiation; Proposed Rule." Federal Register, Vol. 51, 1092, 9 January.

APPENDIX A

1986 March

INTERPRETATION AND PRESENTATION OF RISK ASSESSMENT RESULTS FOR USE IN MAKING RISK MANAGEMENT DECISIONS

Prof. M. Granger Morgan
Department of Engineering
and Public Policy
Carnegie Mellon University
Pittsburgh, PA 15213

tel: 412+268-2672

The author is head of the Department of Engineering and Public Policy and also holds academic appointments in the Department of Electrical and Computer Engineering and the School of Urban and Public Affairs at Carnegie Mellon University. He currently serves on a number of risk related advisory committees for the U.S. Government including the Clean Air Science Advisory Committee of the U.S.E.P.A. He is a member of the Board of the International Society for Risk Analysis, and is Chairman of a Committee on Uranium Mill Tailings of the U.S. National Academy of Sciences.

Quantitative probabilistic risk assessments are all well and good...but what does one do with one of the things once it is done? In this paper I address three issues related to this question:

1. What role should quantitative probabilistic risk assessment play in the overall process of regulatory risk management decision making?
2. Probabilistic results can be confusing, especially to semi-technical and non-technical people. What techniques are available to help clearly communicate such results?
3. What decision analytic tools and techniques are available that would allow one to directly apply the results of a quantitative probabilistic risk assessment in risk management decision making if one wanted to?

1. The Role of Probabilistic Risk Analysis in Risk Management Decision Making

Most people have a mental model of risk management decision making that looks something like Figure 1. A policy maker who is responsible for making risk management decisions identifies the problem that he or she faces and requests a group of analysts, such as the group at Senes, to undertake an analysis. They set up or "frame" the problem, collect needed data (perhaps including subjective expert judgments about data that are needed but are not available) and build models, do the analysis, perform sensitivity studies[1]

---

[1] How much sensitivity analysis gets done depends on the analyst. While sensitivity analysis can be critically important, unfortunately many analysts view this step simply as a formality.

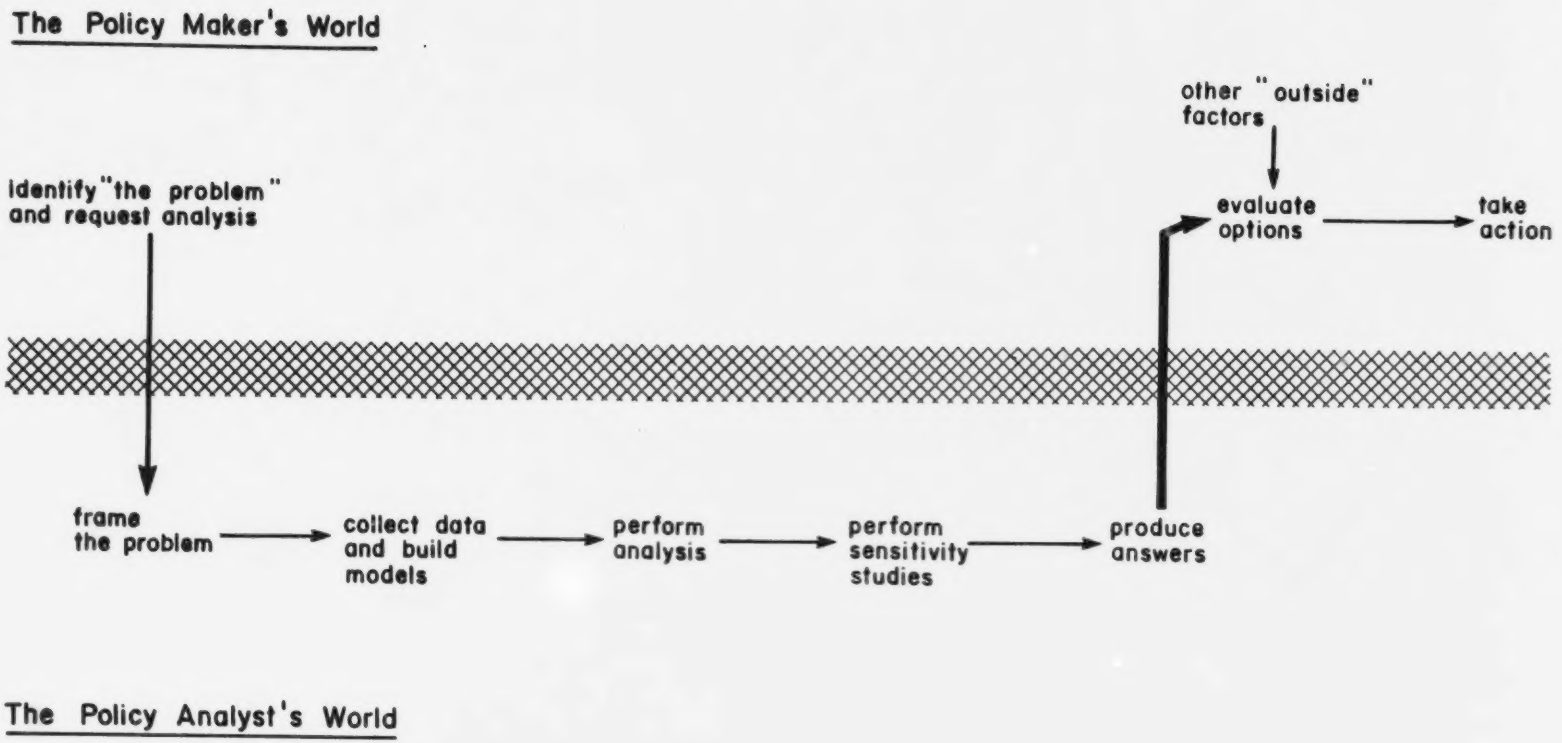
The Policy Maker's World
identify "the problem" and request analysis
other "outside" factors
evaluate options
take action
frame the problem
collect data and build models
perform analysis
perform sensitivity studies
produce answers
The Policy Analyst's World

FIGURE 1. Model of the way in which most people view the role of analysis in the risk management policy process.

and produce a set of findings and "answers" which are communicated back to the decision maker. The decision maker combines these results with various other relevant considerations and reaches some decision.

Under the model of Figure 1 the topic of this paper would be limited to the heavy arrow at the right hand end of Figure 1...that is to the problem of formulating and communicating the results from such a "one pass" policy analysis in a way that maximizes their understandability and usefulness to the policy maker.

While the model represented by Figure 1 probably accurately describes most people's view of the role of analysis, it seriously misunderstands the true value of quantitative analysis in difficult risk and other decision making problems. In such problems, analysis is typically not very good at producing "the answer". It can however often be very good at giving better insight about the problem, and in particular it can help decision makers to refine and improve their framing of the problem and help them to better identify and justify the set of policy options they should be considering.

Figure 2 provides an alternative model of quantitative analysis in risk management decision making which reflects this broader and more iterative role. As in Figure 1, things begin when the policy maker identifies a problem and turns to the analysts for assistance. The analysts develop an initial framing of the problem, collect some initial data and build some simple models in order to do preliminary analysis...with the objective of evaluating the problem framing. That analysis allows the modelers and the decision makers to understand what the important characteristics and constraints of the problem are, what questions it makes sense to ask, what questions are essentially unanswerable, etc. After one or more iterations of this type a revised statement of the problem is jointly developed at which point it then makes sense for the analysts to go on and do more detailed analysis and sensitivity studies. Sometimes in the process of these studies, problems may be uncovered which suggest that one should take another pass at reformulating the overall framing of the problem. Finally, after enough iterations, one will reach the point where one has produced a satisfactory analysis of the final problem. This is then communicated back to the decision maker, *along with a description of the considerations which went into the selection of the final framing of the problem.*

Under the model of Figure 2 the topic of this paper includes the entire process of analyzing the framing of the problem as well as the narrower problems of formulating and communicating the products of the various stages of analysis in a way that maximize their understandability and usefulness to the policy maker.

It is useful to consider the model of Figure 2 in the specific context of the problem of managing the radon risks posed by uranium mill tailings. The initial problem formulation which Senes set out to work on was to produce a single overall assessment of the risk. As their analysis began to unfold it became increasingly apparent that over the course even the next several hundred years (to say nothing of the next several thousand), a wide variety of things could happen at tailings sites that could significantly effect the risks that are posed. It became clear that to try to anticipate and deal with all of these possible developments today would be both silly and economically and technically infeasible.

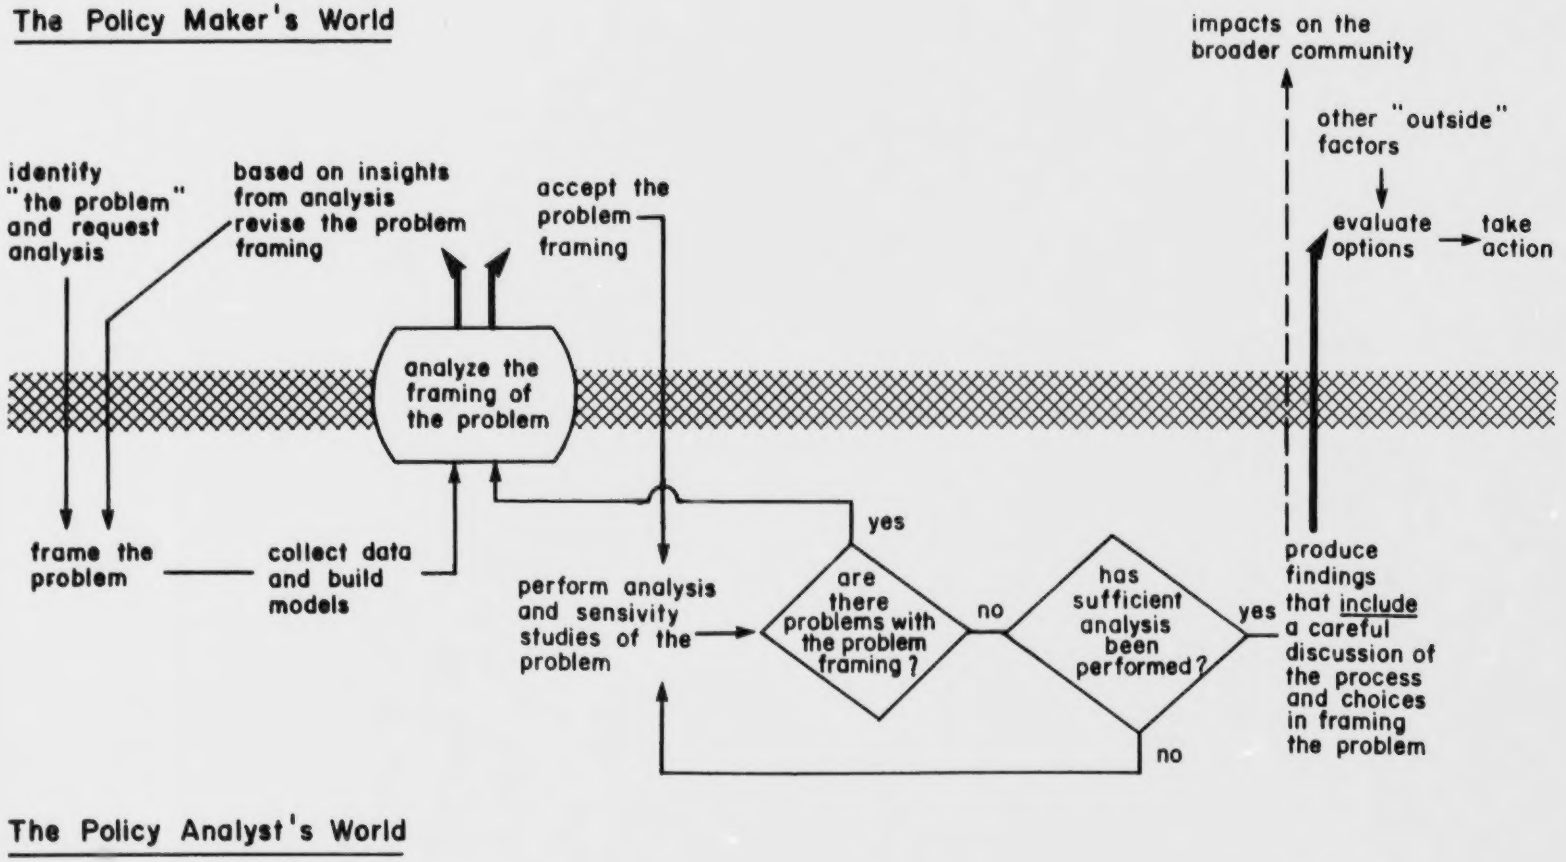
The Policy Maker's World
identify "the problem" and request analysis
based on insights from analysis revise the problem framing
accept the problem framing
analyze the framing of the problem
impacts on the broader community
other "outside" factors
evaluate options
take action
frame the problem
collect data and build models
perform analysis and sensivity studies of the problem
yes
are there problems with the problem framing ?
no
has sufficient analysis been performed ?
yes
no
produce findings that include a careful discussion of the process and choices in framing the problem
The Policy Analyst's World

FIGURE 2. Model of the way in which analysis can make the greatest contribution to the risk management policy process.

In the United States, when the Environmental Protection Agency worked this problem they adopted a model of analysis roughly like that in Figure 1 (1). Without properly iterating on the problem framing to discover the problems that were involved, they adopted a "one shot" solution strategy...cover the tailings piles with 3 meters of earth to impede radon release and then forget them. The analysis that Senes has thus far completed makes it quite clear that such a strategy of abandonment is not an adequate long term solution. Because of the wide variety of geological (e.g. erosion) and ecological (e.g. forest fires) process as well as human activities (e.g. new settlement and land use patterns) which can effect radon releases and exposures from piles, it is clear that any realistic risk management strategy will have to involve some form of "adaptive look ahead" (2). That is, it will have to involve some very low level program of on going attention (e.g. occasional inspections) which can lead to active intervention to correct things like eroding covers when and if it becomes necessary in the future. Once an initial close out has been completed the level of necessary on going attention looks like it could be quite low...the question of how low would be an appropriate question for the next round of analysis. The costs and manpower needs for such a program look like they would be quite modest...the question of who should be responsible, what they should do, how much it would cost and how to pay for it are also appropriate questions for the next round of analysis.

To keep my example simple I have discussed only the radon dimension of the problem. Similar questions should be examined for the ground and surface water dimensions and for the issue of inadvertent use of tailings materials in construction (e.g. back fill around basements). When these questions are looked at in this light I believe they too will lead to an "adaptive look ahead" rather than a "one shot" risk management solution.

My main point is quite simple...the primary value of analytical results of the sort that Senes has been producing is often not the specific numbers or curves themselves, but rather the insights they give about the problem, and the questions that they identify that should be asked and can be answered. In this context the sophisticated uncertainty analysis which Senes has been developing is very important. It makes clear, in a way that is often not possible in more conventional analysis, the limitations to what we can know today and forces us to deal with the resulting ambiguities. If the U.S.E.P.A. had adopted a similar approach they might have come to recognize that an "adaptive look ahead" strategy of the sort discussed above is the only strategy that is compatible with the uncertainty and physical realities of the problem.

Thus it seems to me that the most important thing that Senes and the Canadian National Uranium Mill Tailings Program can do in trying to "interpret and present" the results of risk assessments of the sort that Senes has demonstrated is to focus on the broad issue of whether the risk management problem being addressed has been properly framed, and, if not, what alternative framings seem worth considering. This doesn't take fancy graphical techniques, probability distributions or similar splashy things. It takes careful articulate arguments, and a willingness to observe that "the emperor has no clothes" when the analysis shows this to be the case.

2. Communication of Probabilistic Results

Even if one shares the view of William Hogan that "the goal of quantitative modeling is insight not numbers" (3) one must still display the result of quantitative risk analysis in a way which makes them easy for people to look at and understand. This problem breaks down into two subproblems, first the communication of probabilistic results to technical and semi-technical people who have some training in probability and second the communication of results to semi-technical and non-technical people who have no training or formal experience with probability.

Communicating to those experienced with probability(4).

Most quantitative probabilistic risk assessments involve a significant number of uncertain quantities and model forms. Perhaps the most serious problem one faces in understanding and communicating the results of such studies is the problem of dealing with this high dimensionality. With one or two dimensions one can produce plots of probability density functions or cumulative distributions. When it becomes important to look at a problem across more than one or two uncertain variables, one is faced with the problem of abstracting these dimensions by "projecting" them in some way onto a two or possibly three dimensional display.

We will consider some examples of such projections, but first let us consider alternative one and two dimensional displays. The conventional displays for a single uncertain variable are the probability density function

(PDF) and the cumulative distribution function (CDF), Figure 3. While these two displays contain the same information (the CDF is simply the integral of the PDF) they emphasize quite different things. In general PDFs are best for communicating:

- the relative probability of specific values,
- the most likely value (i.e. the mode) (see Figure 4),
- the shape of the distribution (skewness, kurtosis, etc) (see Figure 4),
- small changes in probability density (see Figure 5).

In contrast, the CDF is best for communicating:

- fractiles, including the median,
- the probability of intervals, including confidence intervals,
- stochastic dominance (see Figure 6)
- mixed, continuous and discrete distributions (see Figure 7).

Since the CDF and PDF are better at communicating different features of the same distribution, it is best, when space allows, to publish both together, one above the other, with the same horizontal scale as illustrated above in Figure 3.

Since the density function is the derivative of the cumulative distribution, it is a much more sensitive indicator of variations in probability density. This may be an advantage if small variations are important in the problem of interest. But if, as is often the case, small

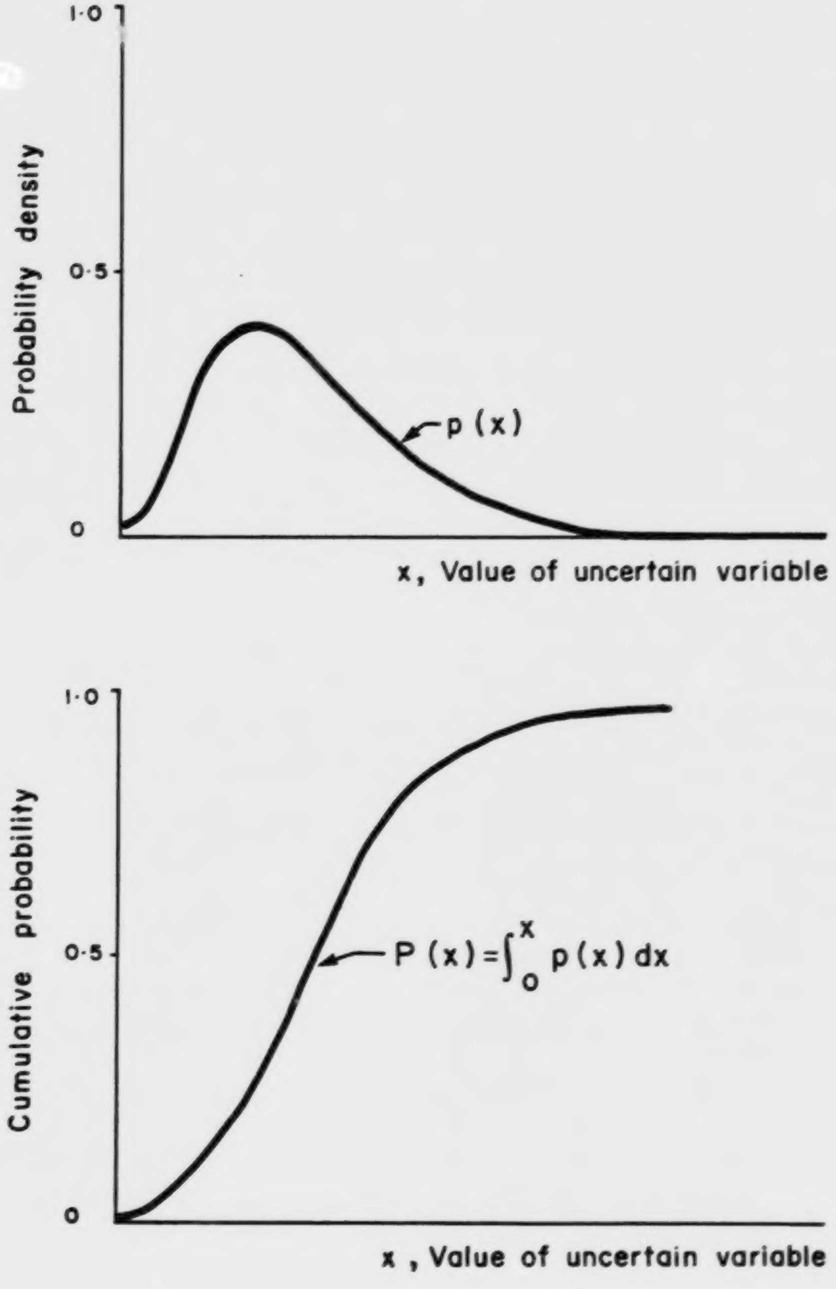

FIGURE 3. Example of a probability density function or PDF (above) and a cumulative distribution function or CDF (below).

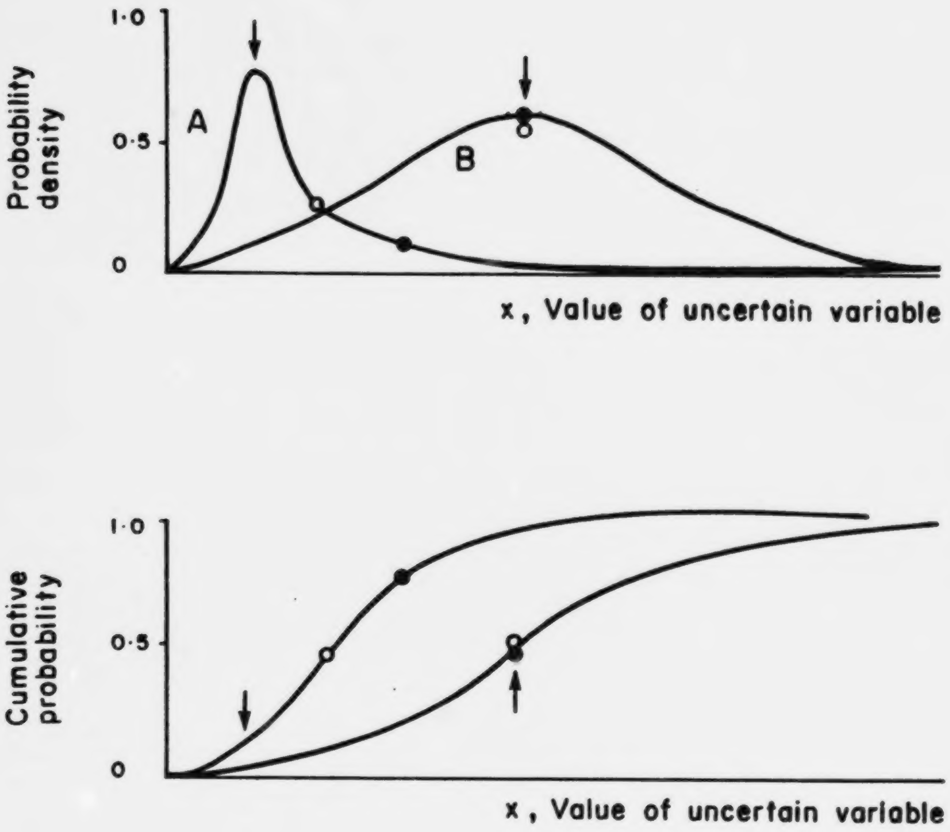

FIGURE 4. The PDF is the more convenient display with which to communicate the "most likely value" or mode, of a distribution (arrows) and to display skewness. Note that distribution A is skewed while distribution B is symetric. Remember that while the "most likely value" is the same as the median (open point) and the expected value (solid points) for a symetric distribution, all three may be different for skewed distribution. Clearly the CDF is best for displaying the median. Neither is superior for displaying expected value.

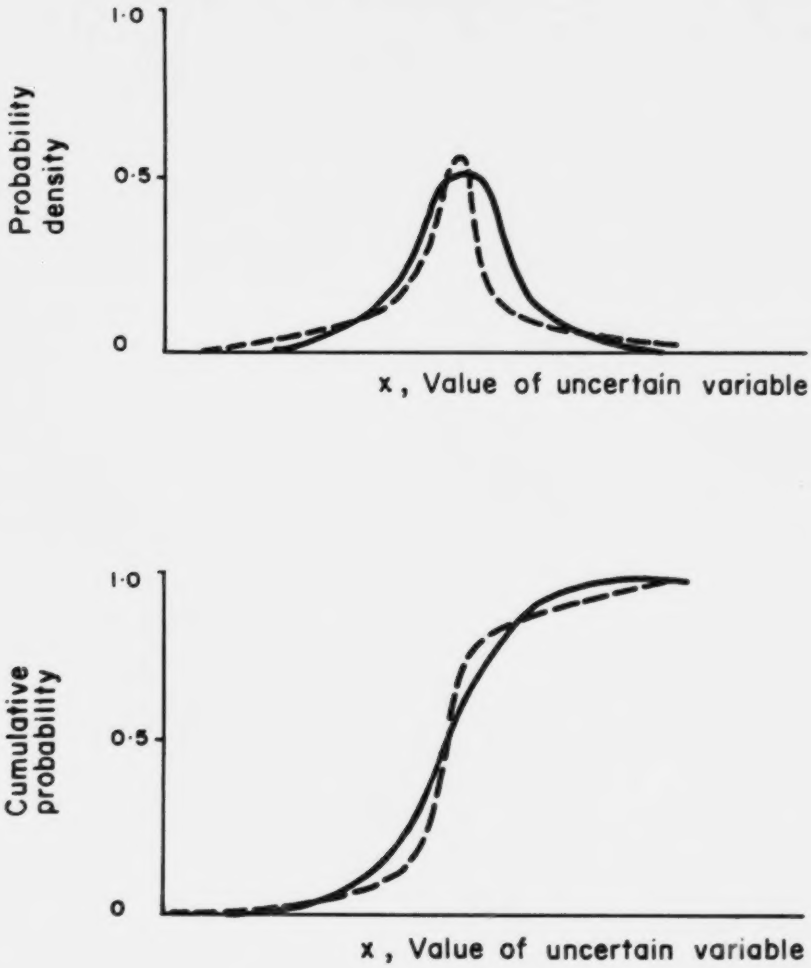

FIGURE 5. The PDF is a more convenient display for illustrating small changes in probability density.

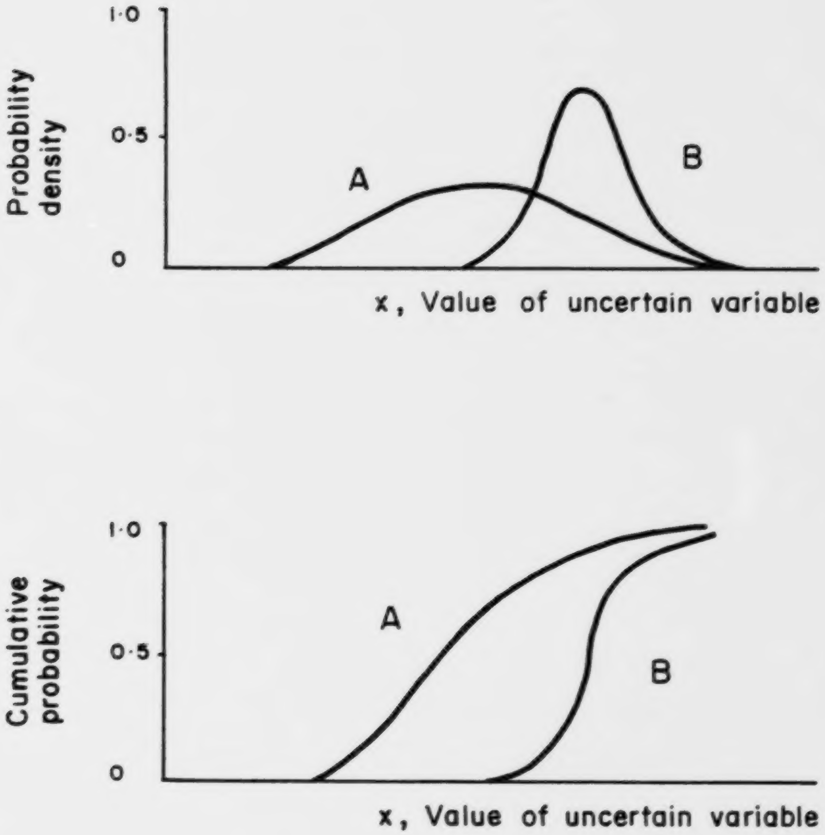

FIGURE 6. The CDF more clearly displays stochastic dominance. In this example distribution B is stochastically dominant over distribution A.

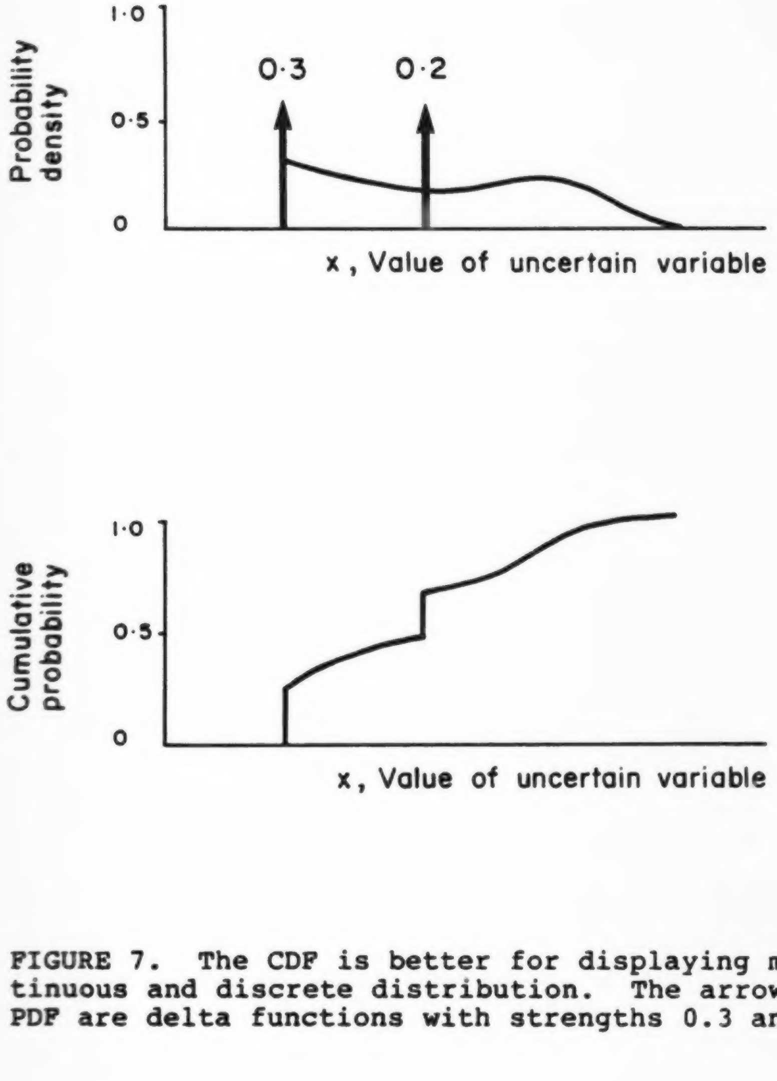

FIGURE 7. The CDF is better for displaying mixed continuous and discrete distribution. The arrows in the PDF are delta functions with strengths 0.3 and 0.2.

variations are not important and the results have been estimated with a stochastic simulation which involved a modest number of model runs, the PDF may display variation which is just random sampling noise as illustrated in Figure 8. Since the convergence of such an estimate to the true distribution goes roughly as the inverse square root of the number of model runs, removing this noise through additional simulation can often be very expensive, and may not be worth while if the process is basically cosmetic...that is it provides no new useful information. In such cases some smoothing of the PDF may be justified.

In addition to the CDF and PDF, various other techniques such as the Tukey "box plot" have been developed for presenting some or all of the same information. We defer a discussion of these to the section on communication to people who are not experienced with probability.

The simplest way in which to deal with two dimensions is to discretize one of them, that is to display a series of one dimensional PDFs or CDFs in the first variable, conditional on the value of the second variable. An example is provided in Figure 9. Similar things can be done with Tukey box plots or the other display procedures discussed in the next section.

An alternative, of course, is to display a perspective view of the full three dimensional surface of the joint CDF or PDF with the two uncertain variables displayed on the two horizontal axes and cumulative probability or probability density displayed on the vertical axis. While this display can be helpful for providing a general sense of the results, it can become awkward if one wants to scale actual values from the plots. In some cases there may also be problems of finding a viewing angle which adequately

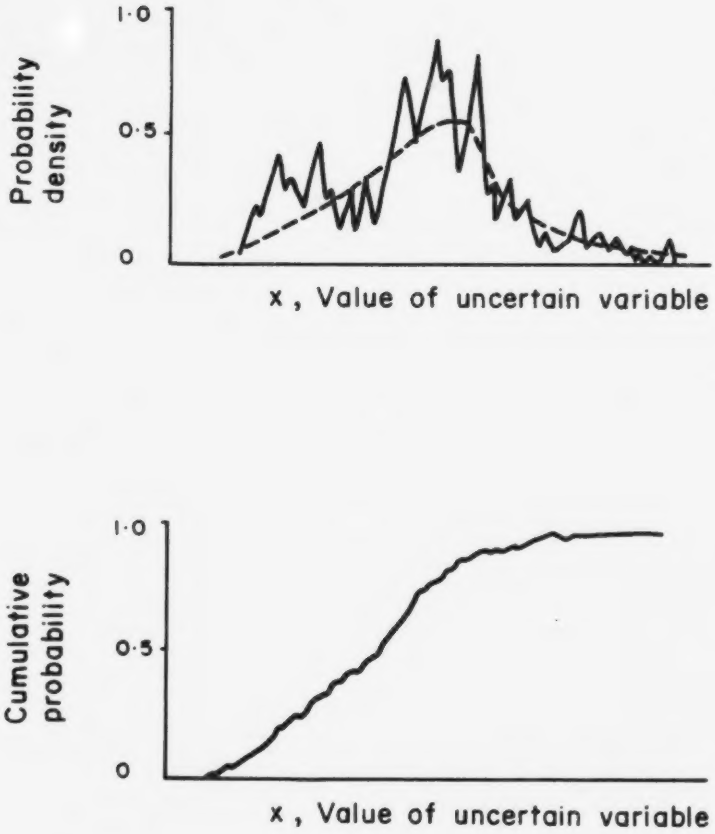

FIGURE 8: Because the CDF is the integral of the PDF [$P(x) = \int_0^x p(x)dx$] it is less sensitive to small variations resulting from random sampling noise in a stochastic simulation.

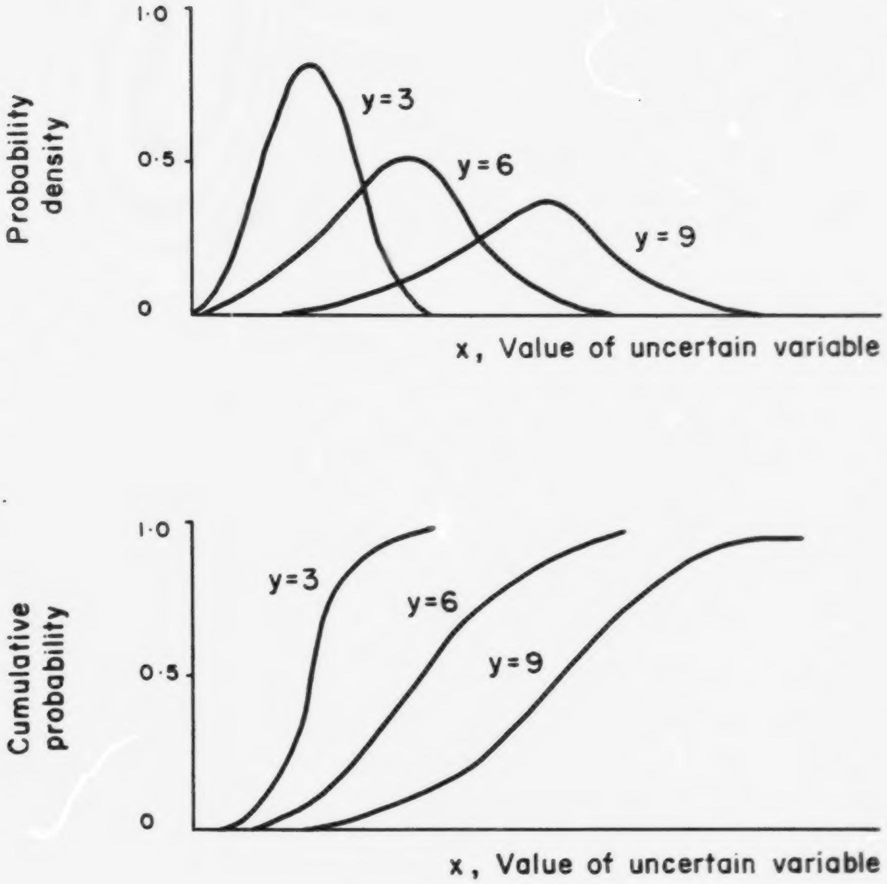

FIGURE 9. Example of treatment of two dimensions (x and y) by treating one as continuous (x) and discretizing the second (y).

conveys the surface.

In the more general case, when one has a large number of uncertain quantities in a model, or even a number of alternative possible model forms, there is clearly no way in which all of the uncertainties can be represented in a single plot.

Faced with this problem of high dimensionality one can describe two alternative extremes. The first, which we can call the *all combinations* strategy involves discretizing all of the distributions (e.g. into high, medium and low) and then running the model for all possible combinations. The typical result is mass of totally indigestible tabular displays. The opposite extreme, that we will label *naive decision analysis*[2] treats all the variables in their full probabilistic form, adds the decision maker's values and attitudes toward risk in the form of one or more (possibly multi-attribute) utility functions and uses the standard techniques of Bayesian decision analysis to integrate across everything and produce a single "answer" which is the policy action with the highest expected net present value or utility. In contrast to the *all combinations* approach which drowned us in so much information about the structure of the problem that we couldn't make heads or tails of it, this approach gives us no information at all about the structure of the problem, it only tells us what to do. It provides no explanation of why this is the best choice, and no insight about how robust this conclusion may be. Clearly the objective must be to strike some appropriate balance between these two extremes.

---

[2] I use the term naive because no good practicing decision analyst would actually adopt the approach I describe.

What one basically needs to do is identify the set of uncertain quantities that are most important and then develop some "projection" from the multi-dimensional model space into a lower dimensional display space. The process is shown schematically in Figure 10. If there is an explicit decision variable in the problem clearly this needs to be explicitly displayed. If there are other key value assumptions, such as the rate of investment to prevent mortality (i.e. "value of life") they also need to be explicitly displayed. Finally, if it is feasible, it is often useful to illustrate the conditional dependency on the value of key uncertain technical parameters.

While the results must be displayed on a two dimensional surface, ingenuity often allows one to display significantly more than two dimensions. Let me illustrate with some examples drawn from some work my colleagues and I did several years ago on possible health impacts from sulfur air pollution(5,6). Figure 11 shows a perspective plot of distributions of sulfur dioxide and sulfate aerosol as a function of flight time in a power plant plume. The plot displays three continuous dimensions: x as concentration, y as probability density, and z as the "index variable" time. Together the two plots also display a forth discrete index parameter, the chemical species. Thus a projection that involves four dimensions is provided. Figure 12 also depicts a projection that displays four dimensions. The plot shows the sulfur mass balance in a power plant plume. The three axes show the partitioning of the sulfur, the contours show discretized contours of a probability density function. In this case the constraint that the sum of the sulfur must add to a constant has been employed to help facilitate the display. Finally Figure 13 combines seven of

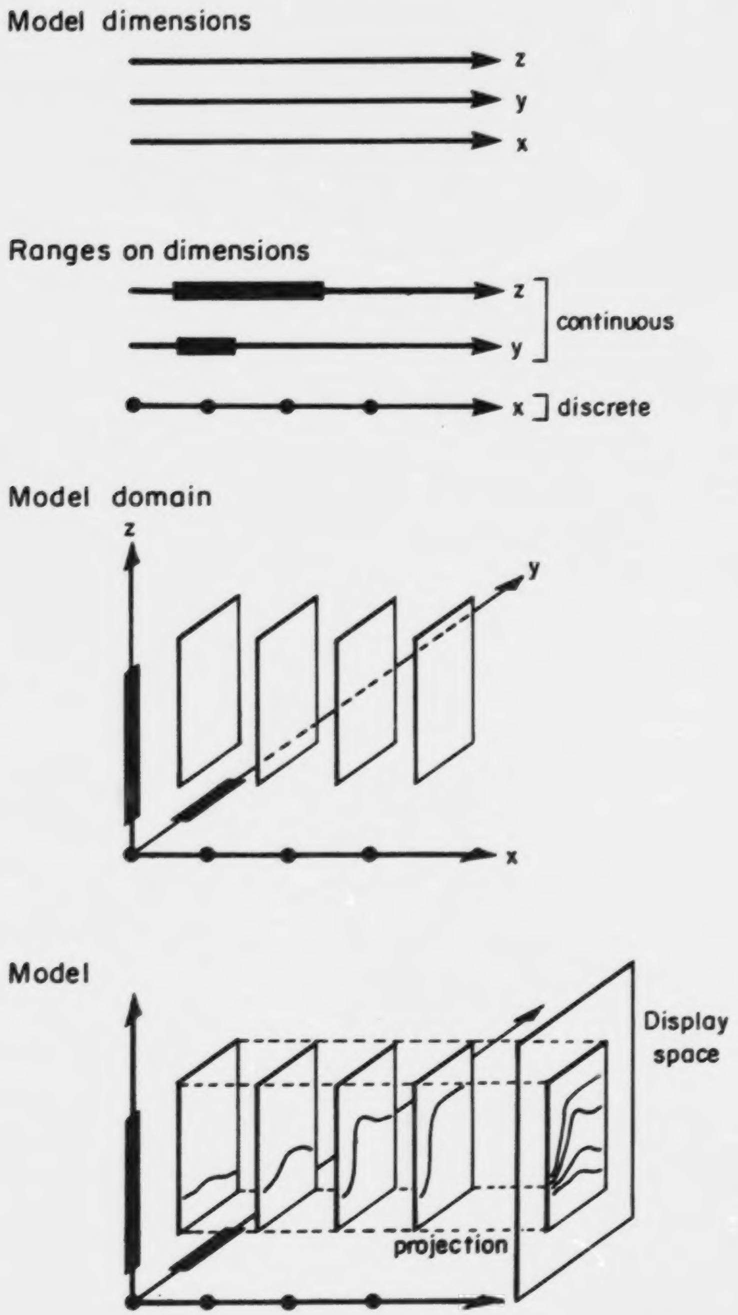
Model dimensions
z
y
x
Ranges on dimensions
z
y
continuous
x
discrete
Model domain
z
y
x
Model
Display space
projection

FIGURE 10. The process of representing results that involve several dimensions in a two dimensional display space can be thought of as a process of "projection".

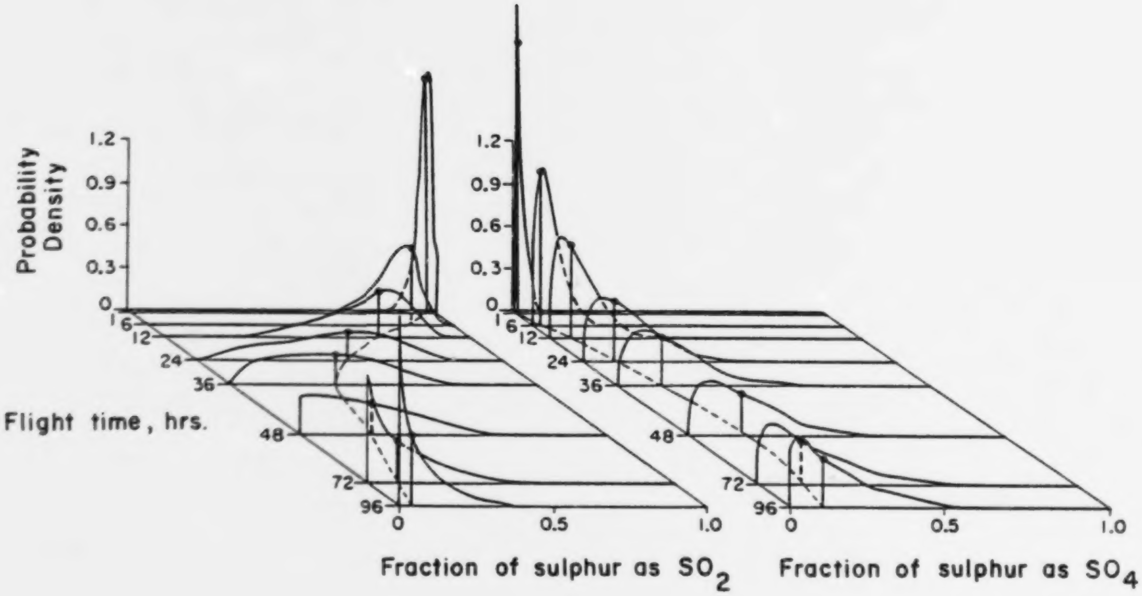

FIGURE 11. Probability density functions for the total amount of sulfur emitted from the plant which is present as sulfur dioxide (left) and sulfate aerosol (right) as a function of flight time. Reproduced from Morgan (5).

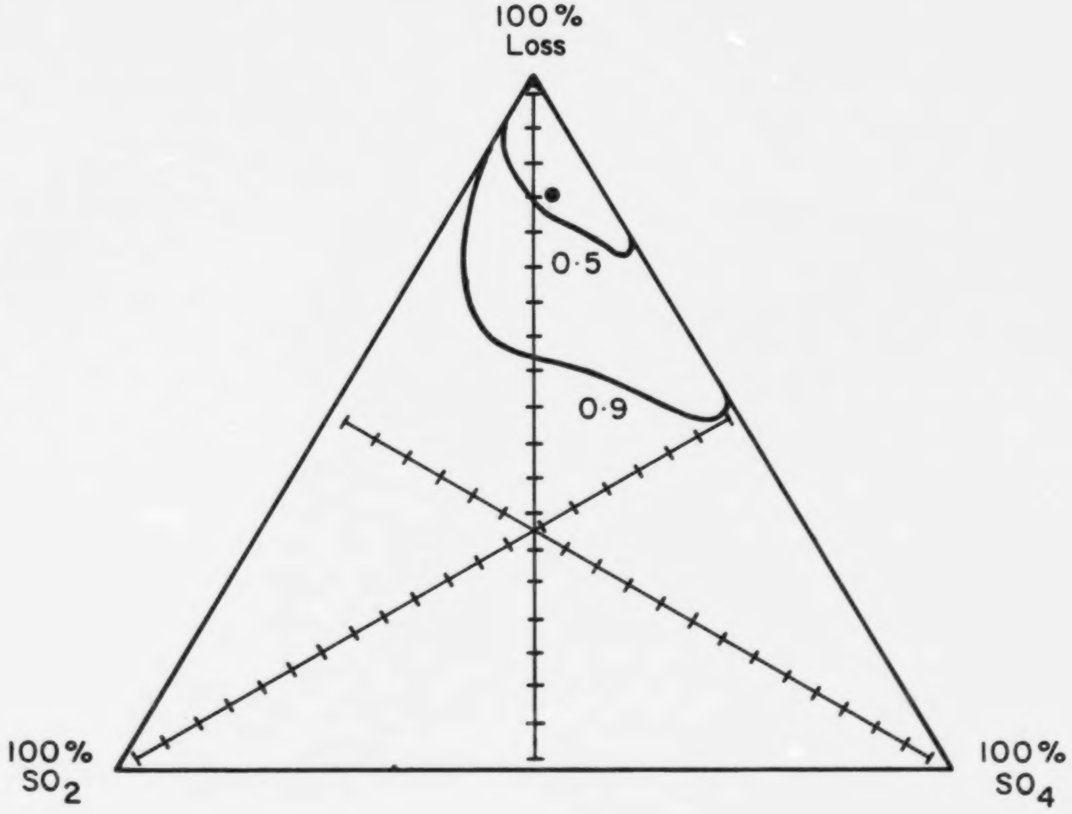

FIGURE 12. Plot of sulfur mass balance in the plume of a coal-fired power plant, after Morgan (5,6).

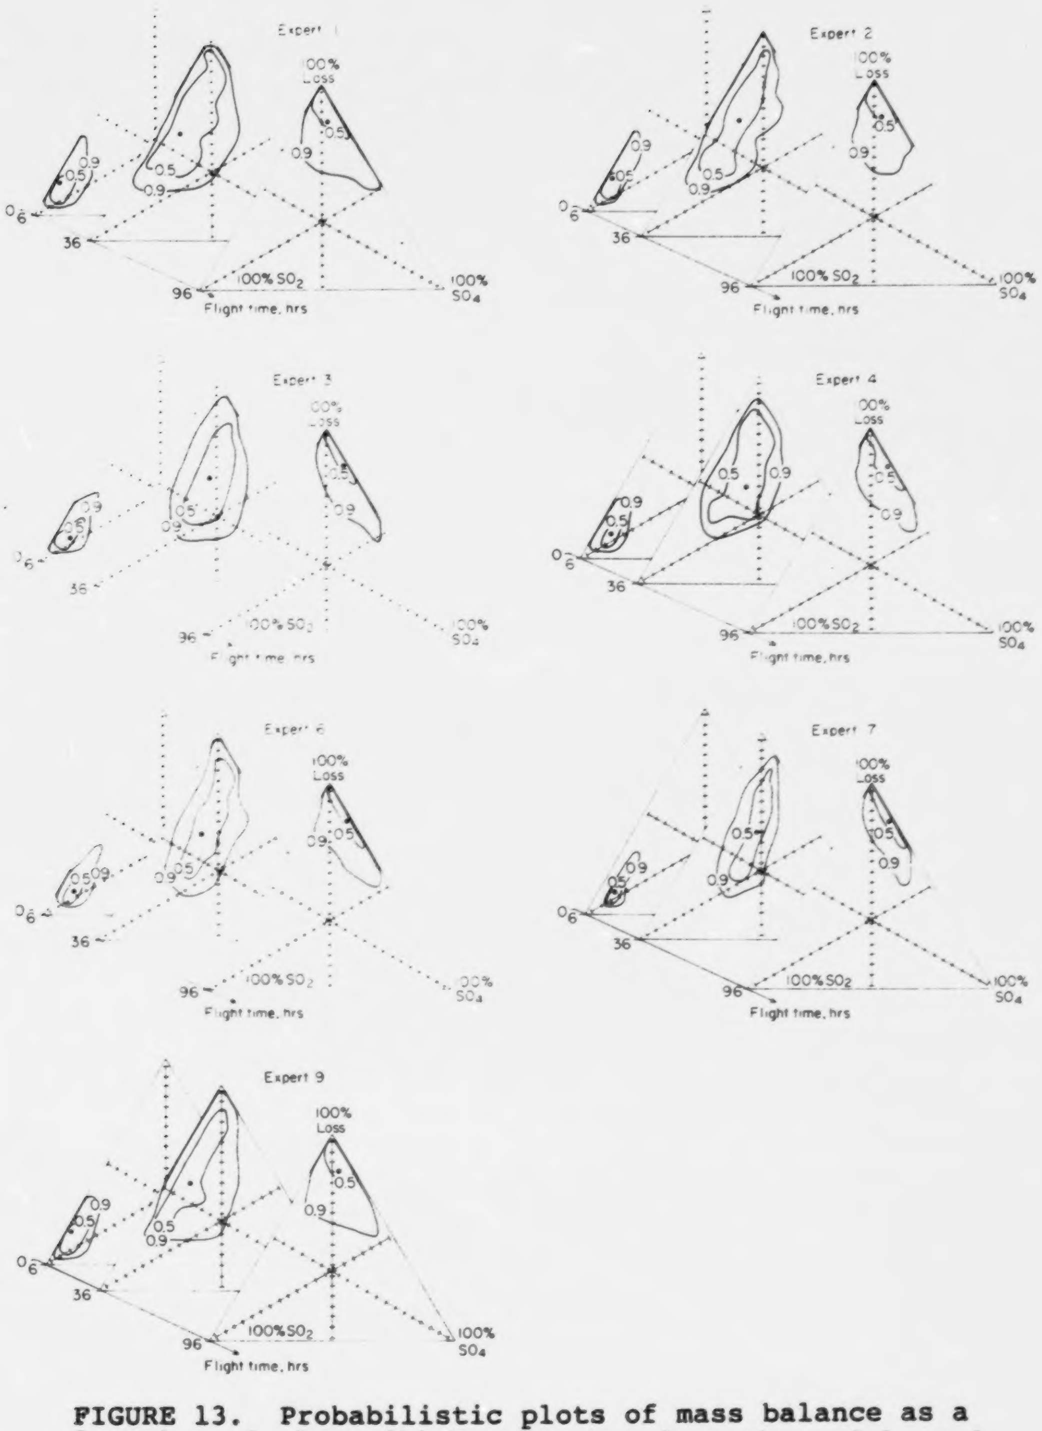

FIGURE 13. Probabilistic plots of mass balance as a function of plume flight time based on the models and subjective probability distributions provided by each of seven atmospheric science experts. Inner contours report the 50% confidence region. Outer contours report the 90% confidence region, after Morgan (5).

these plots (corresponding to the results for each of seven experts) at each of three flight times to produce a projection that displays three continuous and three discrete parameters for a total of six.

With some care in construction, displays of this sort can be very helpful for people who have a reasonable understanding of probability. Such highly complex displays are considerably less appropriate for people who are not familiar with the formal notion of probability. We turn next to the problem of displaying information for these people.

Communicating to those not experienced with probability.

While many people have never had any formal exposure to the ideas of probability, almost everyone has some basic understanding of uncertainty. Further, if the ideas are introduced in familiar terms such as in the context of probabilistic precipitation forecasts, or point spreads on horse races, it turns out most people rather quickly catch on. Once they get the basic idea most people find the idea of displaying estimates of uncertainty about things that they intuitively know must be uncertain to be very appealing.[3]

However, there are limitations to how accurately one can convey probabilistic results to people who have not been formally schooled in the basic ideas of probability. In order to begin to understand these limitations, together with Harold Ibrekk, I have been running a series of experiments that involve the communication of single uncertain variables,

---

3. At a luncheon address several years ago for the annual public meeting of the Illinois State Environmental Agency I used such an approach to present the results on sulfur air pollution described above to a non-technical audience with very considerable success (7).

such as snow or flood depth to non-technical people. We have completed a pilot study and are in the midst of data analysis on a larger study. Unfortunately only tentative findings based on the rather small sample from our pilot are available at this time. The full results should be available later this spring.

We evaluated nine alternative displays in this study. These included conventional PDFs and CDFs, Figures 14 and 15; a discretized bar chart representation of the density function, Figure 16; a Tukey box plot that displayed mean, quartiles and the full range of the data, Figure 17; a point estimate of the mean along with a horizontal line that displays the 95% confidence interval, Figure 18; two alternative bar representations of the density function, Figures 19 and 20; a representation of probability density constructed by halving the density function and then displaying it along with its mirror image in the x axis, Figure 21; finally we used a pie chart representation as illustrated in Figure 22.

The results obtained to date are not strong enough to report in detail. However several findings do appear reasonably strong and, if taken as still preliminary, deserve mention. We observed that there was a bimodal response to CDFs. Some people understood them and interpreted them correctly, However, a significant number clearly do not understand them at all. For example when asked to select the "most likely" value this group selected the maximum possible value. While a non-technical explanation of the various displays, including the CDF, improved performance, it did not eliminate the misinterpretations. This preliminary finding suggests that CDFs alone may be an unreliable way to communicate information about uncertainty to semi-technical and non-technical audiences.

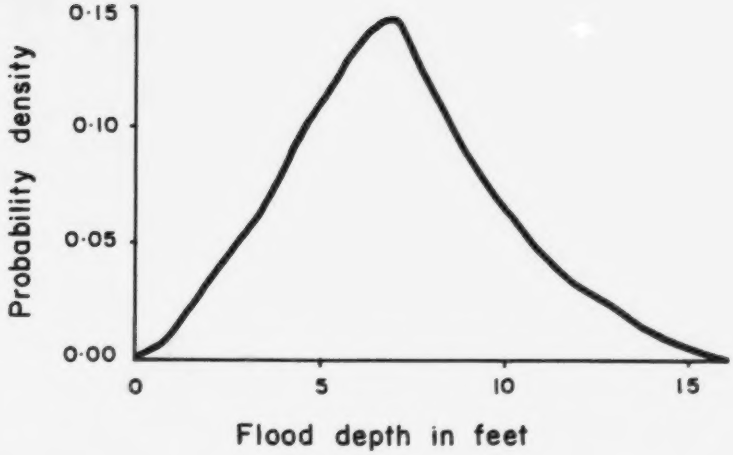

FIGURE 14. Example of probability density function used in uncertainty communication experiment.

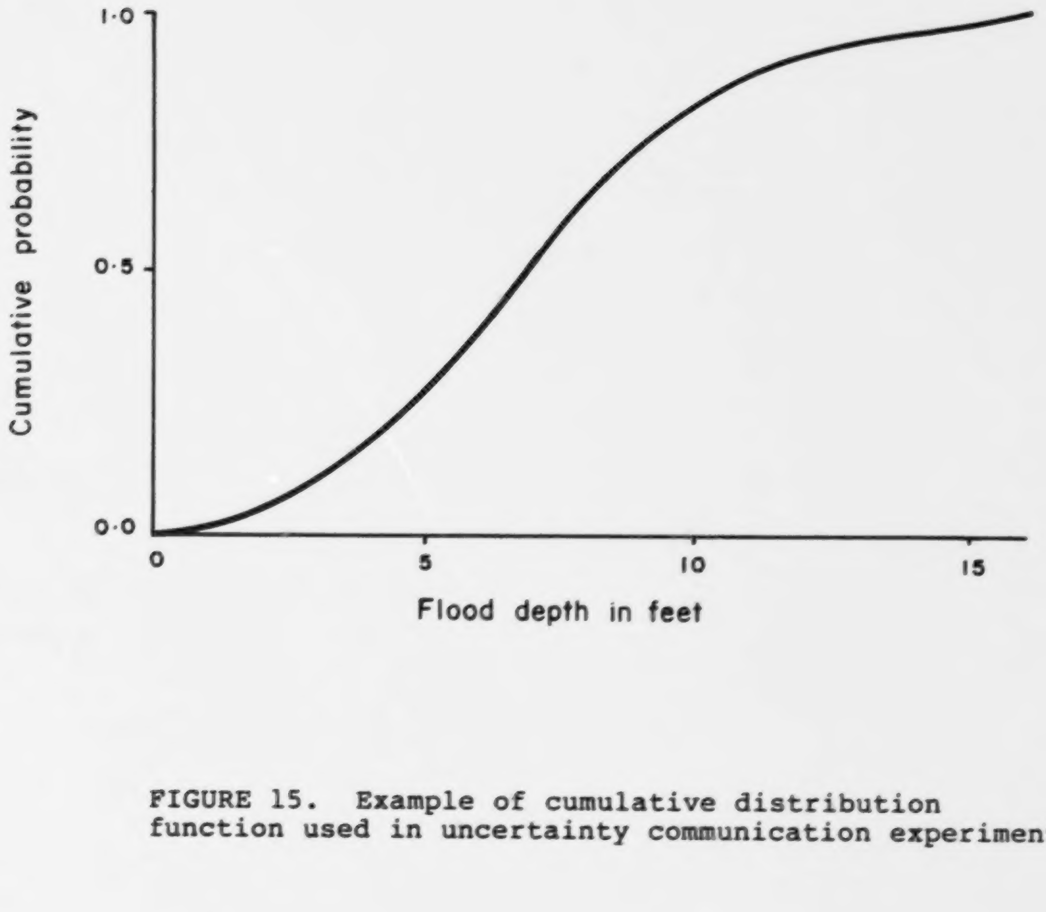

FIGURE 15. Example of cumulative distribution function used in uncertainty communication experiment.

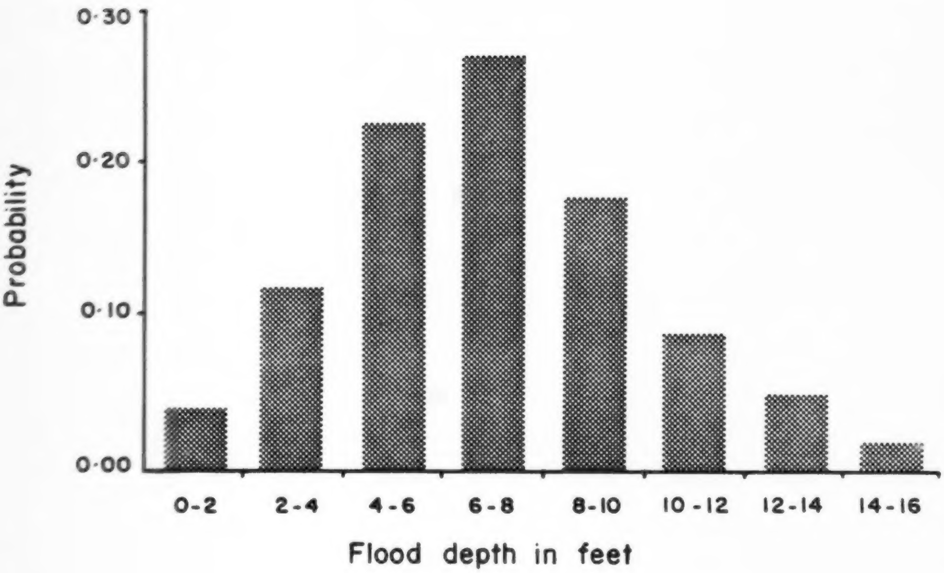

FIGURE 16. Example of discretized probability distribution used in uncertainty communication experiment.

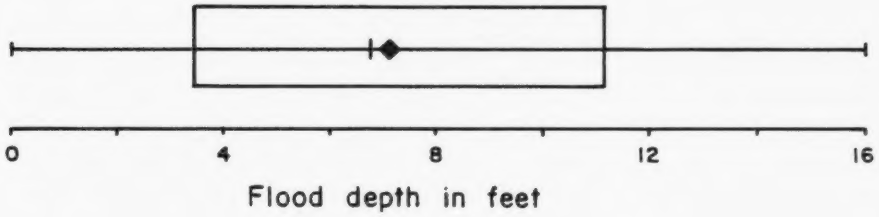

FIGURE 17. Example of modified Tukey box plot used in uncertainty communication experiment. The box spans 75% of the probability, the bar spans the entire range of possible values, the vertical mark is the median and the dot is the expected value.

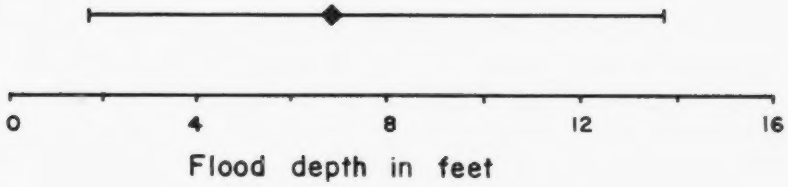

FIGURE 18. Example of point and range display used in uncertainty communication experiment. The dot is the expected value. The bar gives a 95% confidence interval.

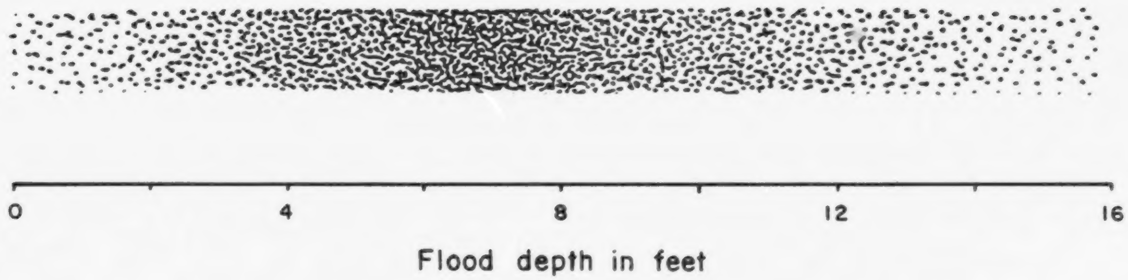

FIGURE 19. Modified probability density display used in uncertainty communication experiment.

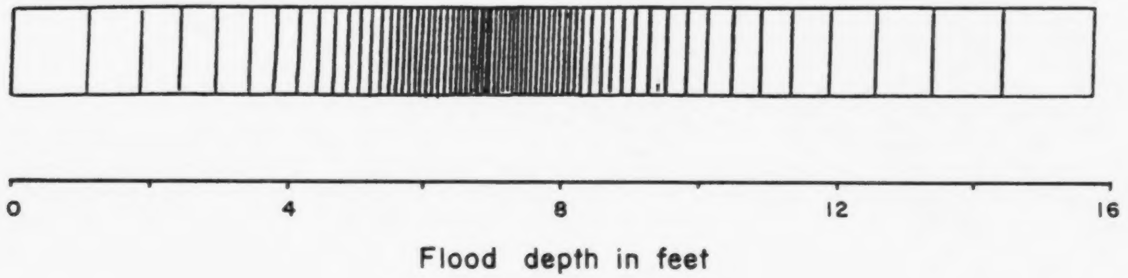

FIGURE 20. Modified probability density display used in uncertainty communication experiment.

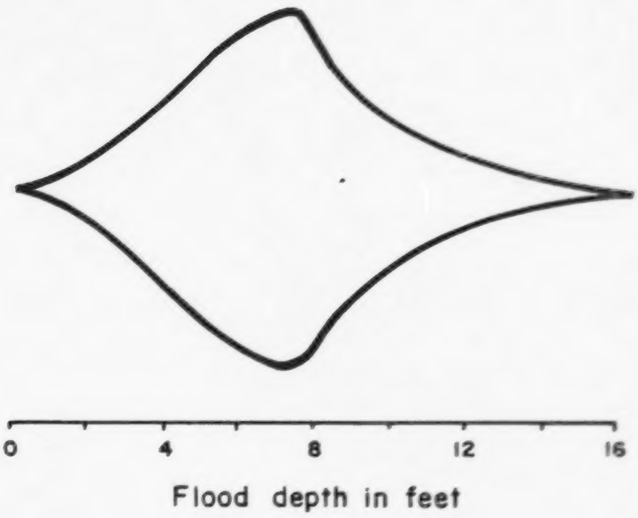

FIGURE 21. Modified probability density display used in uncertainty communication experiment. In pilot studies this particular display performed especially well.

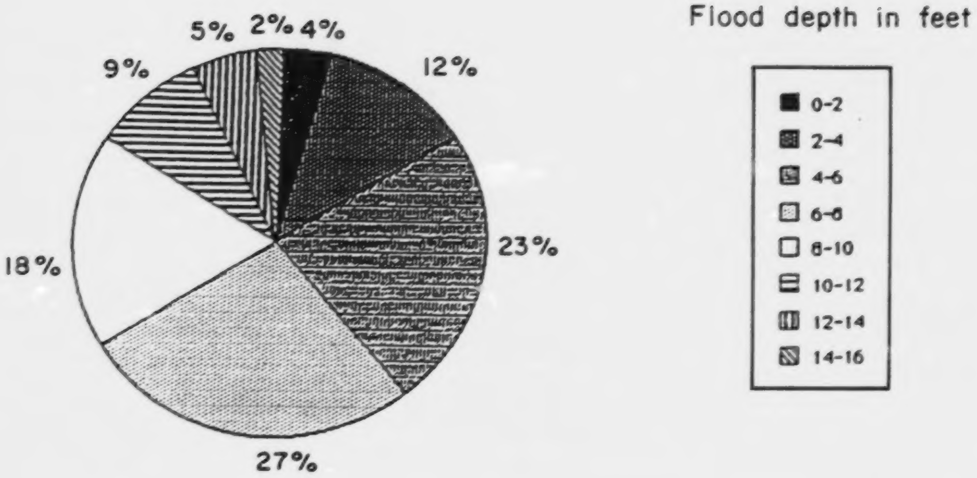

FIGURE 22. Pie chart probability display used in uncertainty communication experiment.

Our non-technical subjects appeared to do better (in estimates of mean, probability of a range, and probability of a result above a specified level) with the displays shown in Figures 16, 17, 18, 19, 21 and 22 than with Figures 14, 15, and 20.

While the finding is not very firm, our subjects appeared to do the best overall job interpreting, and deducing information from, the representation of Figure 21. It is not yet clear which, if any if these preliminary findings will be confirmed in our larger study, now underway.

3. Use of Decision Analytic Tools

There are two broad ways in which probabilistic results of the sort represented in Figure 23 and 24 can be used in support of risk management decision making:

1. Qualitatively, to give the decision maker a general "feel" for the situation and some basis for making a choice among ontions without actually applying any explicit quantitative decision criteria, or
2. Quantitatively, as part of a variety of possible alternative decision analytic approaches.

While the bulk of this section will be devoted to a discussion of the second of these approaches, one should not understate the value and importance of the first. Often problems become sufficiently complex, or so many important factors in the decision process can not reasonably be quantified, that formal decision analytic approaches become unworkable, or at

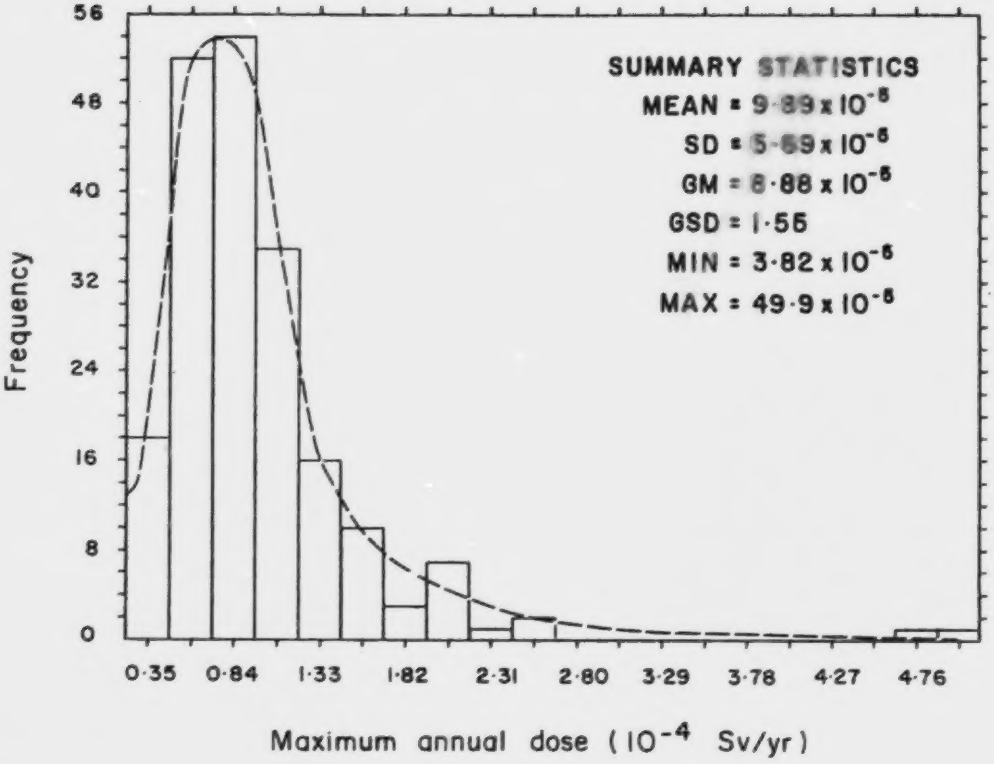

FIGURE 23. Example of the kind of probabilistic output used in the NUTP risk assessment. Readers should think of the dashed curve as roughly equivalent to the continuous curves in Figure 25, on the left hand side of Figure 26, etc.

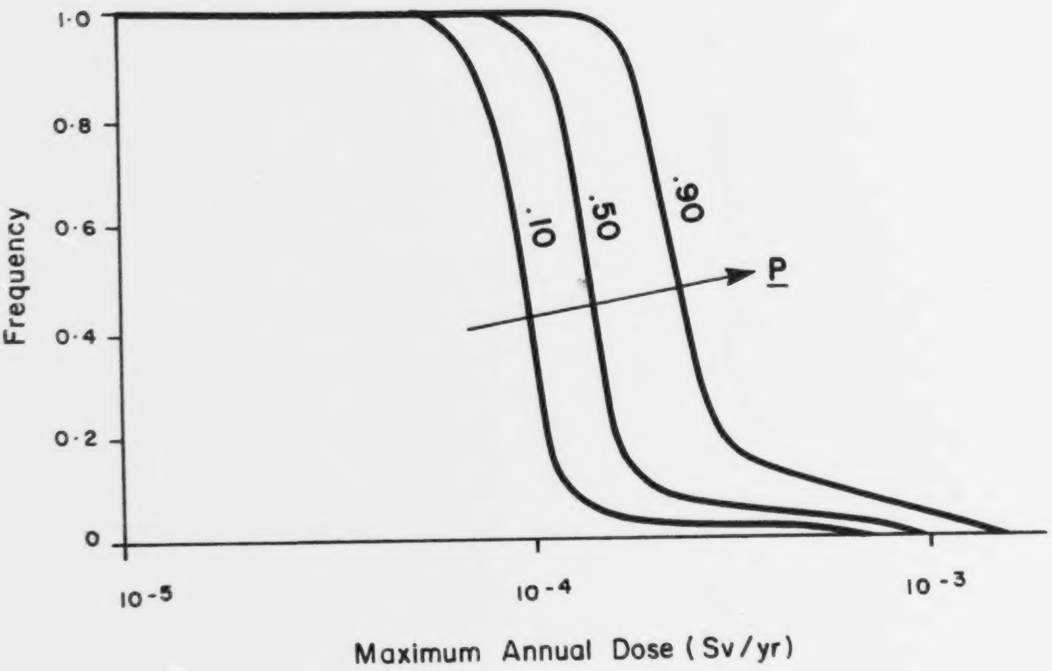

FIGURE 24. Example of probabilistic output which introduces "second order" uncertainty, that is uncertainty about the frequency distribution of future events.

least fail to provide decision makers with insight and a defensible basis for their decisions. Sometimes in such circumstances looking carefully at the results of all the available assessments and then just "doing what seems right" may be as or more defensible a decision procedure as any formal quantitative process. However this is true only if one has thought carefully about the basis on which the decision is being made and the factors that are important. Formal decision analysis is helpful in forcing such thought, and thus may be useful even in situations in which the final decision is made "by the seat of the pants".

Faced with making a choice between a variety of alternative risk management actions that might be taken to control the potential risks posed by something like uranium mill tailing piles, one must have some basis for evaluating alternatives. In the real world the factors which enter in such a decision may be quite complex and may include a variety of political and other considerations. For the purposes of this discussion we will take a narrower prospective. Risk management objectives which might be adopted in making such decisions include:

1. Absolutely limit exposures (or levels of effects) below some specified level,
2. Limit probability of exposures (or levels of effects) below some specified level(s),
3. Maximize expected net social benefit (or minimize expected net social loss),
4. Apply various alternative single attribute objective functions such as

minimizing the maximum possible loss,

5. Apply a multi attribute utility formulation and apply some multi-attribute objective function.

Combinations of these objectives, or other objectives not listed here, are also possible, but this set is sufficient for our discussion.

Absolute limits: At first glance the simplest formulation of the problems would appear to be in terms of an absolute limit that must not be exceeded under any circumstances. For example, in Figure 25 one would choose option A over option B because A meets the limit L, and B does not.

A limit can be stated in terms of a variety of measures. It might for example be stated in terms of maximum annual dose in which case a presentation such as that shown in Figure 23 or 25 would be an adequate basis for evaluation. It might on the other hand be stated in terms of a maximum lifetime risk of cancer in which case one would have to disaggregate the sources of exposure that go into Figure 25 into a number of contributing sources, each of which would then be multiplied by an appropriate risk factor or health effect function and the result summed in order to get total cancer risk. A hypothetical result is shown in Figure 26.

There is also the issues of whether the limit is stated in absolute or incremental terms. For example if the limit is stated in terms of maximum annual dose, is it a limit in total dose from all sources including "natural" background or is it a limit on incremental dose above "natural" background? A hypothetical illustration of the difference is shown graphically in Figure 27.

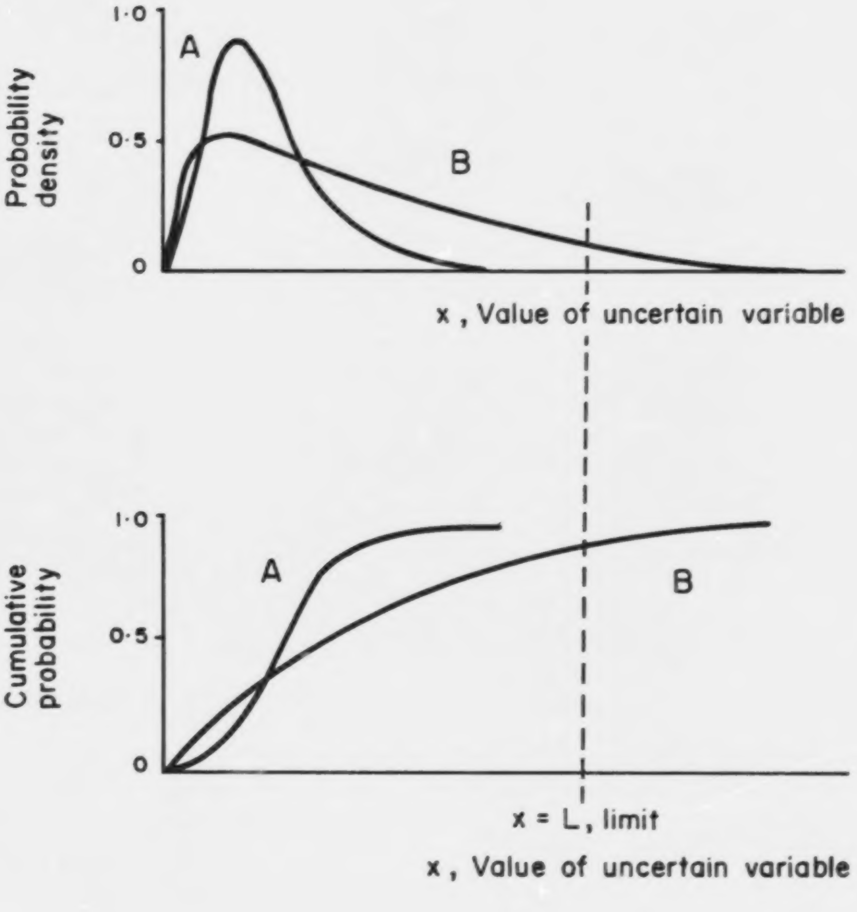

FIGURE 25. If one is controlling a risk through the application of an absolute limit (L) on the variable (x) (i.e. the risk management strategy is $x < L$) then control option A would be acceptable and control option B would not be acceptable.

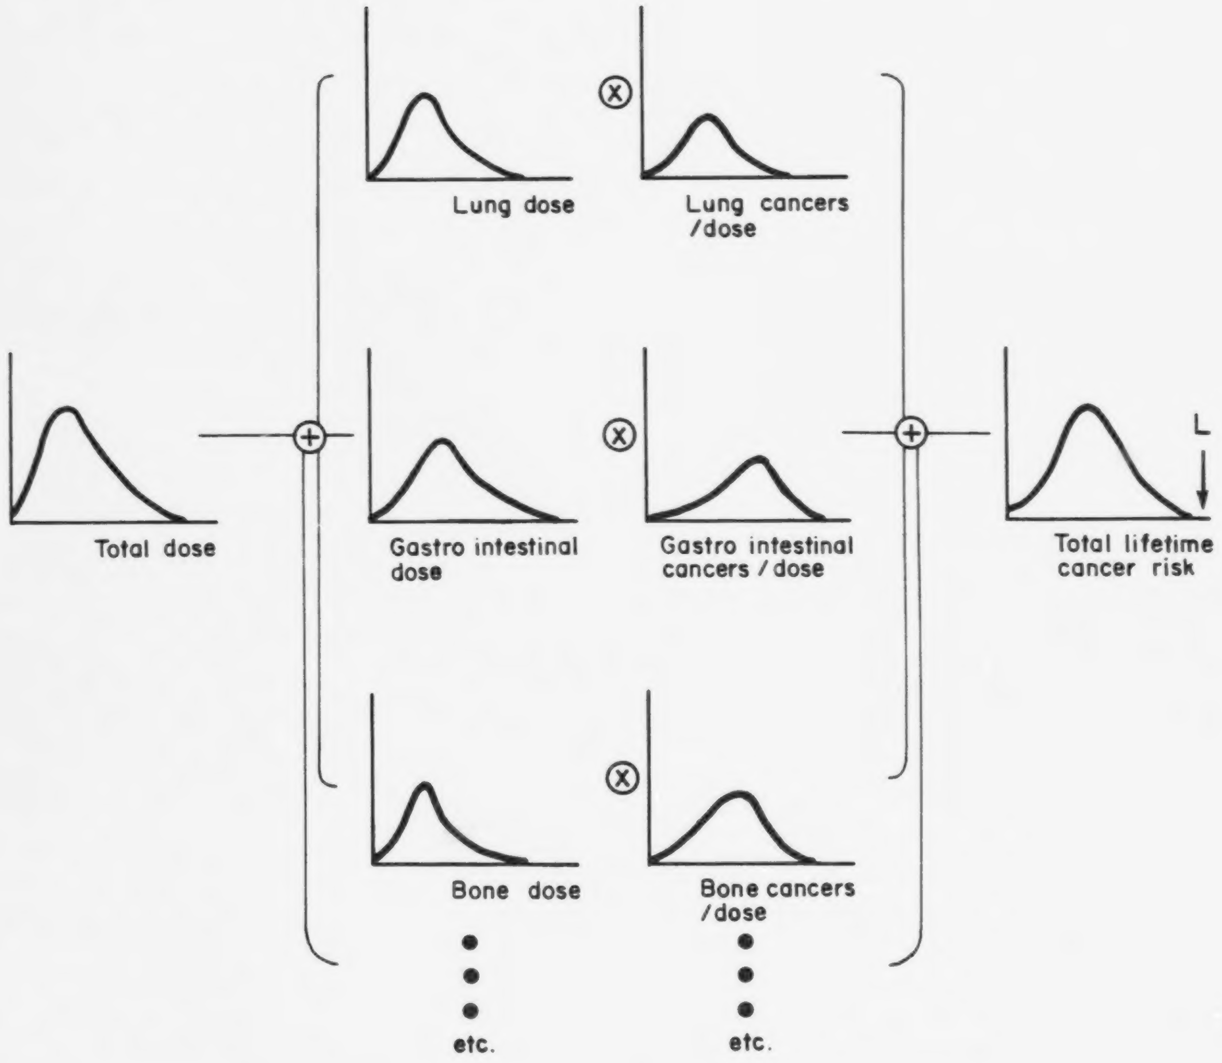
Lung dose
Lung cancers /dose
Total dose
Gastro intestinal dose
Gastro intestinal cancers / dose
L
Total lifetime cancer risk
Bone dose
Bone cancers /dose
etc.
etc.

FIGURE 26. Limits, (L), may be stated in a variety of terms. If (L) were stated as an upper bound on the acceptable incremental lifetime cancer risk for specific types of cancer, then output of the sort displayed in Figure 21 could not be used directly. Instead it would have to be disaggregated into separate organ doses to produce separate cancer risk estimates which could then be re-aggregated into a total risk. (Alternatively, total cancer risk could be estimated directly from total effective dose using ICRP dose factors.) Note that all curves shown in this illustration are hypothetical and do not reflect what the actual distributions would look like.

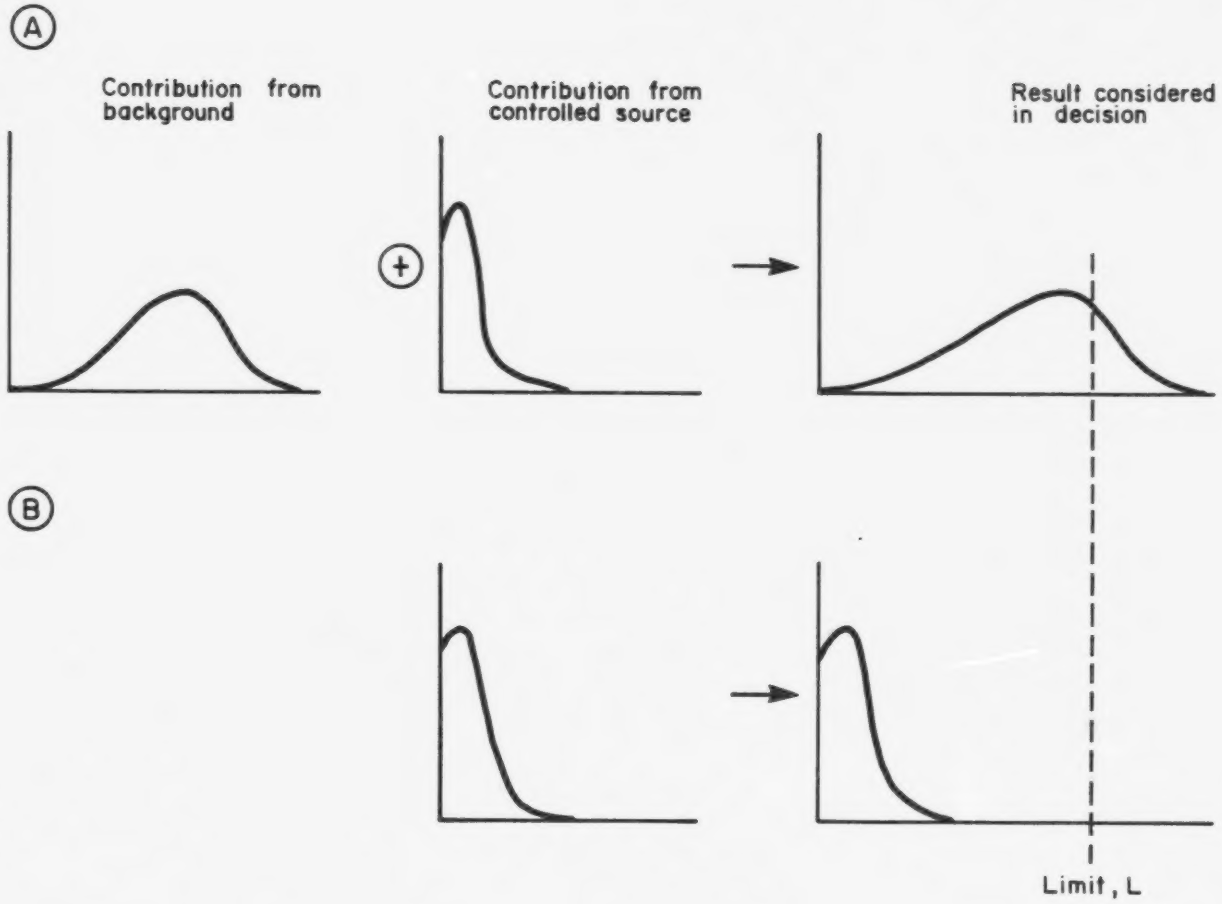
A
Contribution from background
Contribution from controlled source
Result considered in decision
+
B
Limit, L

FIGURE 27. In considering a risk management strategy based on the imposition of an absolute standard (L), one must decide whether or not to include contributions (to risk, dose, etc.) from background (A), or to consider only incremental contributions (B). The choice is not always obvious. For example, should lung cancers from the voluntary activity of smoking be included when one is protecting against lung cancer?

Probabilistic limits: A problem that typically arises when one is faced by a situation such as that shown in Figure 25 is that there is a limit to the certainty with which one can specify the high tail of the distribution. That is, there is a very low but nevertheless finite probability of exposures that exceed L. Typically in adopting an engineering strategy for risk management, it is far less expensive to limit the size of the high tail than it is to eliminate it entirely. This gives rise to a different kind of limit, one that is stated as a probability. For example in designing structures like highway embankments it is common to design them to withstand the hundred year flood, but not the thousand year flood. Thus one does not say "design this road embankment so that it will never fail" (a design specification which over long enough periods of geological time is clearly impossible), rather one says "design this embankment do that the annual probability of failure is <0.01".

In principal probabilistic limits could be stated in more complex terms than as limits, for example as a full CDF which must stochastically dominate the actual distribution, Figure 28. I am unaware of any actual applications of such an approach although U.S.E.P.A. appears to be moving this way with some of its ambient air quality standards. There are probably cognitive and other limitations on how sophisticated probabilistic limits of this type can become and remain practical.

Maximum expected social benefit: While deterministic and probabilistic limits have their appeal, they can also present problems because they do not (at least directly) incorporate a consideration of cost. In some cases, in which society believes that the appropriate decision rule is a "rights based" approach to risk management in which considerations of control cost are not relevant, this approach is fine. For example, there are some risks that the

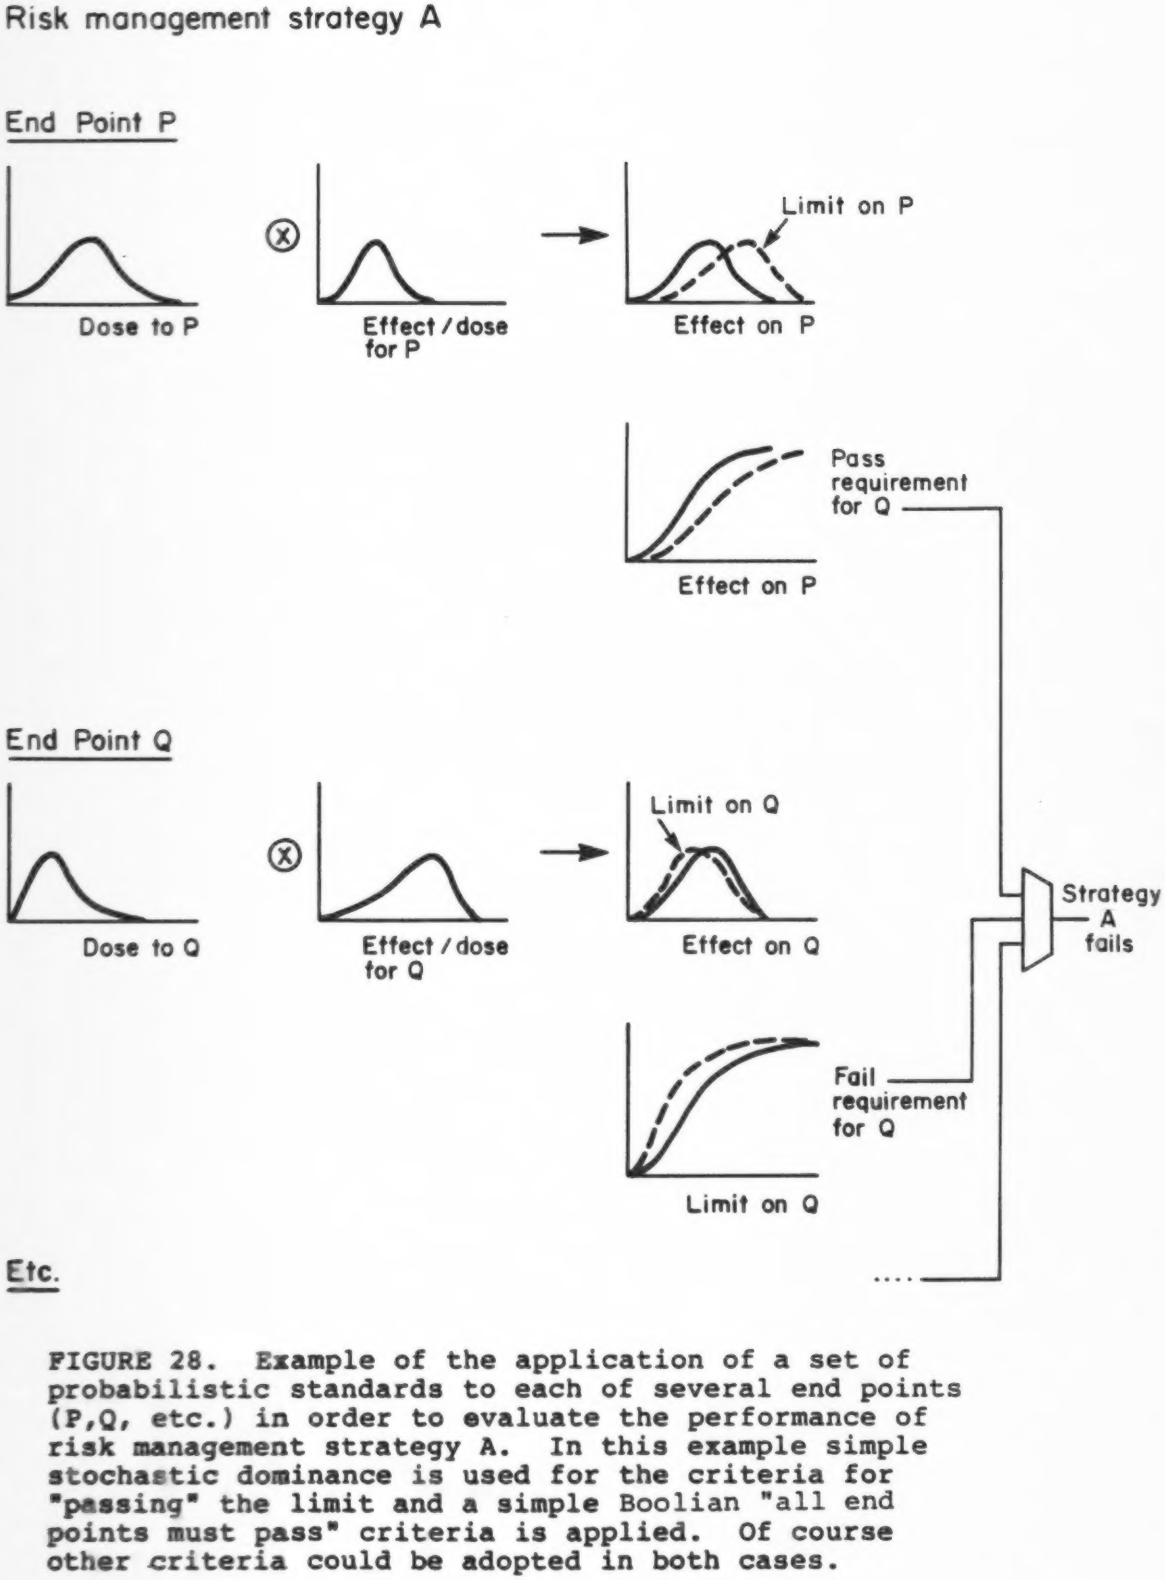
Risk management strategy A
End Point P
Dose to P
Effect / dose for P
Limit on P
Effect on P
Pass requirement for Q
Effect on P
End Point Q
Dose to Q
Effect / dose for Q
Limit on Q
Effect on Q
Strategy A fails
Fail requirement for Q
Limit on Q
Etc.
....

FIGURE 28. Example of the application of a set of probabilistic standards to each of several end points (P,Q, etc.) in order to evaluate the performance of risk management strategy A. In this example simple stochastic dominance is used for the criteria for "passing" the limit and a simple Boolian "all end points must pass" criteria is applied. Of course other criteria could be adopted in both cases.

law says my neighbor can not impose on me independent of how much benefit may accrue to him or society from his doing so. In other cases, in which a utilitarian approach based on maximizing net social benefit is accepted as the appropriate decision rule, a limit-based approach to risk management is simply not feasible.

We illustrate such a case with the hypothetical example in Figure 29. Suppose that there are three alternative risk management strategies, call them A, B, and C. The impacts of these strategies in terms of incremental cancer risk are illustrated in the lefthand set of curves in Figure 29. Each of these also has an associated cost of implementation as shown in the second set of curves in Figure 29. In order to estimate net social cost or benefit one must select a rate at which society is prepared to invest to prevent incremental cancers. Since there is unlikely to be unanimity about the value that this investment rate should take it is typically best to treat it parametrically.[4] The result is a set of distributions in net benefit of the sort shown on the right hand side of Figure 29 in which the distribution in net present value is computed for each of a number of different safety investment rates R. For any given investment rate the option with the lowest expected cost is the preferred option. If the choice of option does not show a strong dependency on the investment rate, then a choice can probably be made with some confidence. If however the choice of preferred control option depends critically on the investment rate across a socially plausible range of such rates, then the decision maker clearly needs to give the problem

---

[4] These rates have been shown to range from several hundred thousand dollars per death averted to tens of millions of dollars per death averted in a variety of risk management programs of the U.S. Government (8).

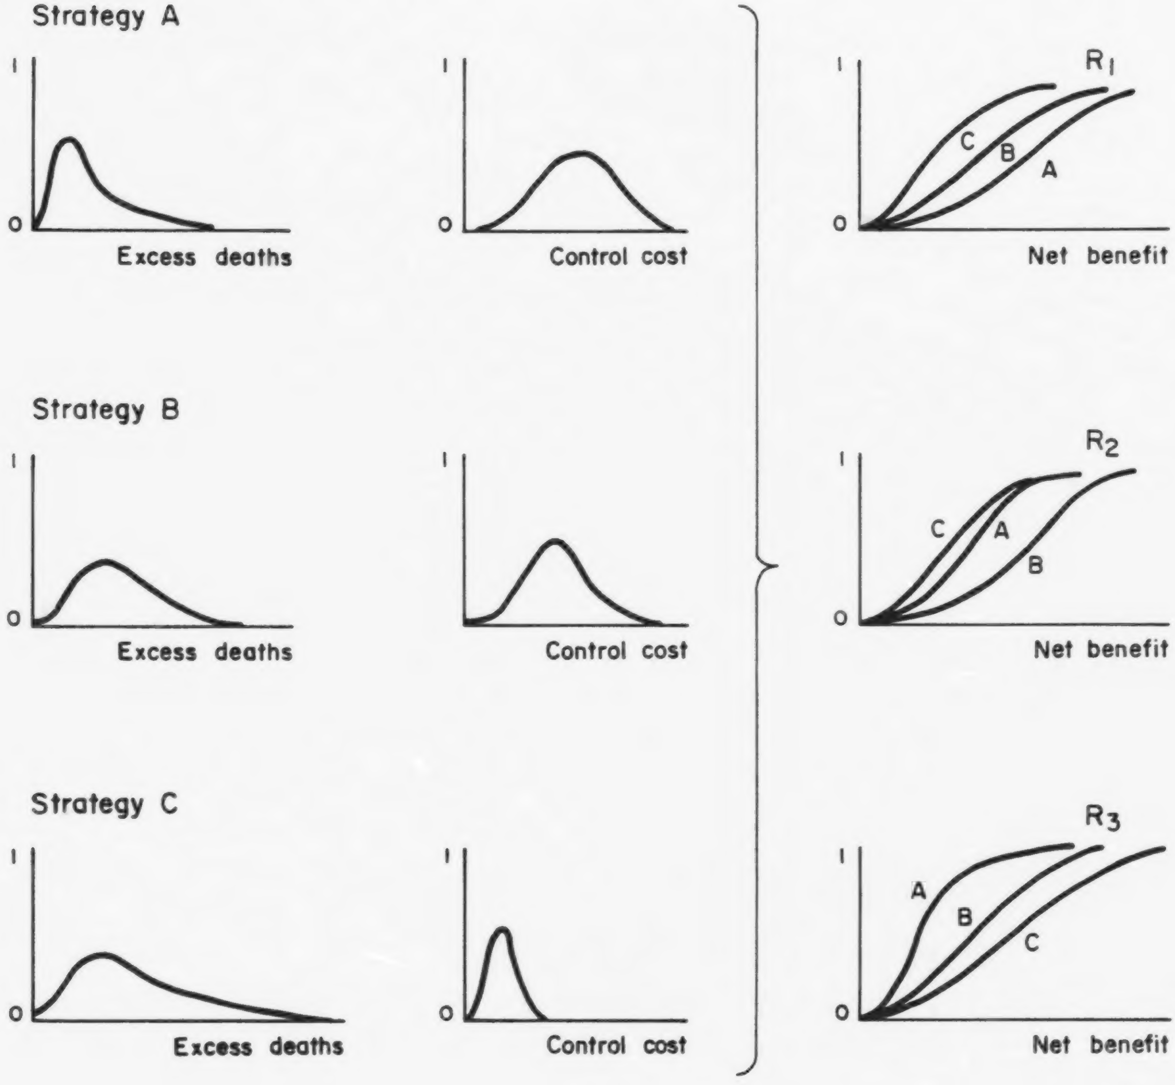

FIGURE 29. Illustration of net benefit calculation which incorporates both the cost of alternative control strategies A, B, C and alternative rates of social investment to prevent mortality, R_1, R_2, R_3. If, as is often the case (8), results are not sensitive to R across a plausible range of R's then the choice of strategy is clear. If, as in this illustration choice of strategy depends critically on R, the decision maker must work to develop new insights or reframe the problem in a way that will allow meaningful discrimination between alternative strategies.

careful additional consideration.

The approach I have outlined with a specific example is known by a number of names such as Bayesian decision theory, optimal statistical decision theory, decision analysis, etc. There is a large technical literature on this subject (9). While my example was presented in terms of dollar costs and benefits, under some circumstances dollars may not provide an adequate measure of value. This can be the case, for example, for an individual or a society which has a great deal, or very little of a particular good or service, or in situations in which the society or individual is risk averse. In such circumstances decision theory employs the concept of "utility" to incorporate the fact that not all dollars or other goods and services are of equal value under all circumstances. The necessary utilities are assessed using fairly standard procedures, and then the problem is worked through in terms of utility rather than dollars. Once the optimal outcome has been selected one can then reconvert to dollars to display the "monetary equivalent" value of the selected strategy. In the context of this paper these are really technical details. Since the ideas are well developed in the literature, I will not provide a tutorial here. Interested readers can find details in the literature (9).

Minimum maximum cost...and similar strategies: While the decision analytic literature argues persuasively for the use of expected net benefit, or, if value is not adequately represented by dollars, expected utility, there are circumstances in which decision makers have good reasons to want to apply other decision criteria even if society as a whole were to assert (for example through a piece of risk management legislation) that maximum expected net benefit is the appropriate decision criterion. One such example is the circumstance in which a public official is prepared to accept a higher overall expected loss (which is nevertheless small enough so that society will tolerate it) in order to avoid a small but finite risk of a larger loss which, while it is very unlikely to occur, the political process would not tolerate if it happened. In the unlikely event that this larger loss did come to pass, the political process would hold the decision maker and his political party responsible. An example is shown in Figure 30.

Of course, viewed from the perspective of the decision maker's preferences rather than those of society as a whole, it may still be possible to reframe this example in terms of maximizing expected net utility.

Multi-attribute utility: One problem with the kinds of approaches I have been discussing above is that for real environmental problems in which the costs and benefits involve a variety of different things like lung cancers, forest ecosystem diversity, water quality, individual freedom, etc. it is typically difficult, and often impossible to reduce everything to a single common metric such as dollars. Further, the amount of money a society (or an individual) is prepared to invest in forest ecosystems, or in preventing lung cancer, is likely to vary depending upon things like how much healthy forest it has, how healthy its population is, and how much disposable income it has.

Multi-attribute utility theory has been developed as a formal decision tool to assist in situations of this sort (10). In principal one could assess a multi-attribute utility function by asking a decision maker to choose between a large number of pair-wise gambles on vectors of outcomes. In practice the number of choices become too large. A large part of the

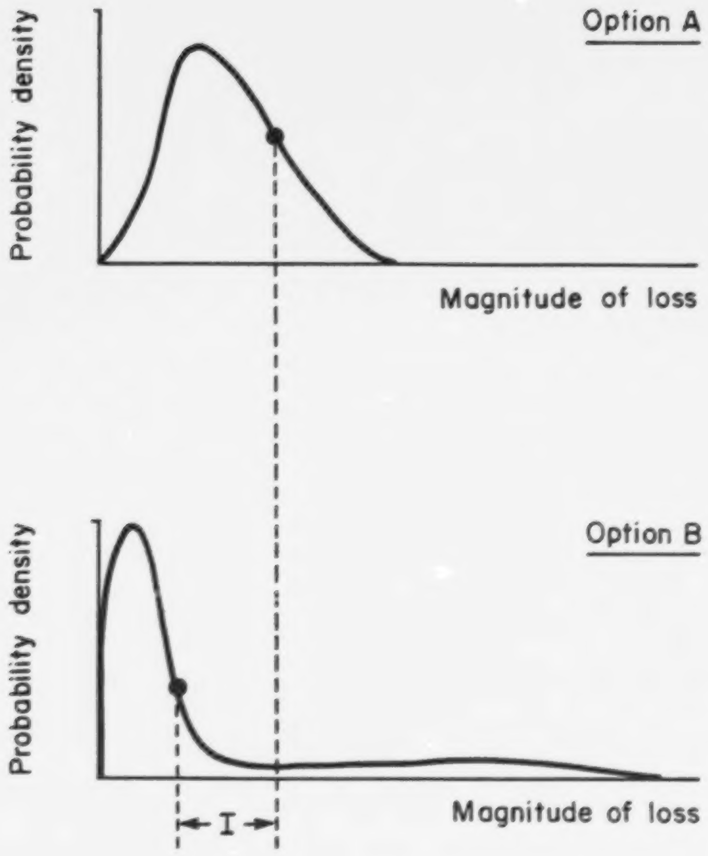

FIGURE 30. Illustration of a situation in which society may indicate a preferance for Option B because it minimizes expected loss (median value, solid point) but a public decision maker might choose Option A with increased expected loss (I) because he or she is applying a "minimize maximum possible loss" criterion, since large losses, though rare, have unacceptable political consequences for the decision maker.

literature in this field is devoted to developing a set of "independence tests" which allow one to ask a much smaller set of tradeoff questions, make some functional assumptions about the shape of the utility function, and then fit it with the results of a small number of additional tradeoff questions. Since from the point of view of the discussion in this paper these are just technical details, I will not discuss them further here. Interested readers can find a detailed discussion in Keeney and Raiffa (10).

Uncertainty about uncertainty: Most of the results that the Senes analyses are producing are in the form of simple PDFs and CDFs. However, in the final round of work they have also begun to produce results such as those in Figure 24, which involve second order uncertainty, or in the language that Kaplan has adopted, uncertainty about the expected frequency of future random events.

In order to make this presentation consistent with the graphical conventions that I have adopted throughout the preceding discussion, Figure 31 replots Figure 24. The specific curves shown in Figure 24 are the output from an analysis, but let us begin our discussion by considering how such a family of curves might be interpreted if they were to be presented as an expert's subjective judgment for use as the input in an analysis. Suppose that such an expert, when interviewed using classical procedures for the elicitation of expert subjective judgment (11), produced a result such as that shown in Figure 32a. Now suppose that I go on to ask this expert questions about how much I can change the curve he has given me and have him still believe that it represents (or more or less represents) his views. For example I might show him a bunch of curves such as those illustrated in Figure 32b and ask him the extent to which he can distinguish between these

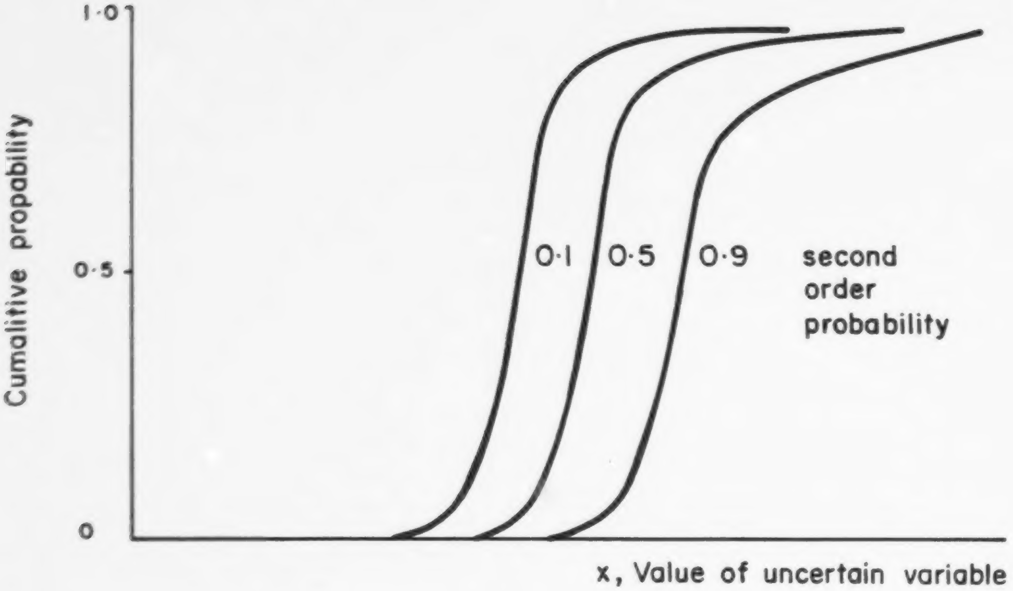

FIGURE 31. Replotting of Figure 24 to make it consistent with the conventions that have been used in this discussion.

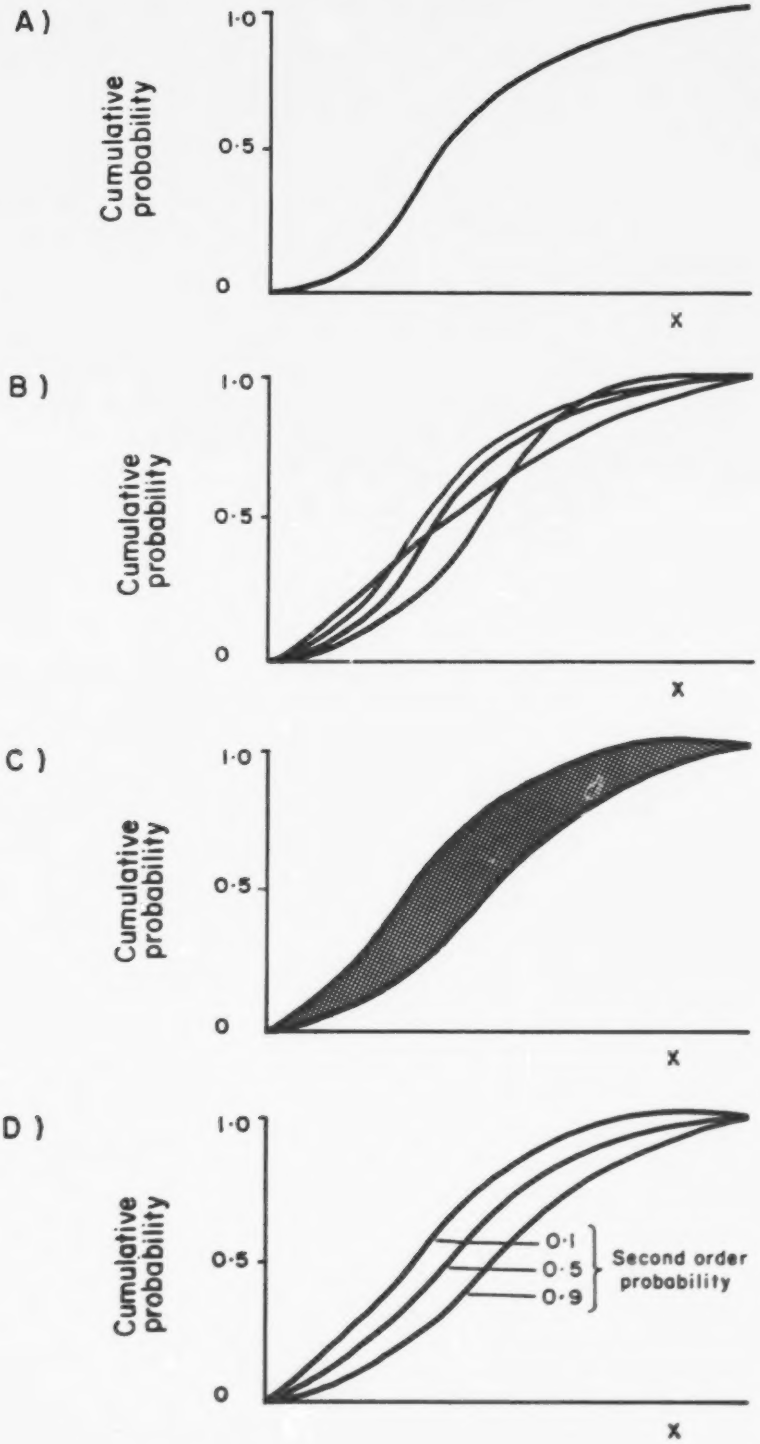

FIGURE 32. Illustration of construction of second order uncertainty in a CDF elicited from an expert (see text for discussion).

and the curve he gave me. If the expert has read the experimental literature on expert calibration (12), he will know that his judgment is probably not reliable enough to distinguish meaningfully between a family of such similar curves. Thus (given a variety of assumptions and some analytical manipulation which for simplicity in this discussion I will gloss over) one might be able to represent his judgment by an envelope of curves such as that shown in Figure 32c or with further effort by a result such as Figure 32d, which is a result in the same form as Figure 31.[5]

There is a second way in which one could imagine constructing a curve such as that represented in Figure 31. Suppose that one has a large number of alternative models of physical reality, each of which produces a slightly different outcome. The result might be a family of curves quite analogous to that shown in Figure 32b. With some appropriate additional judgments, one could then proceed from this to produce a curve like Figure 32d.

Whether it is useful to preserve information about second order uncertainty depends entirely on what one intends to do with the result. Preserving the second order uncertainty presents no problems if one adopts an absolute limit approach. In the case of an absolute limit one simply insists that all the mess of curves in Figure 31 lie below the limit...in this case the details of the shapes of the curves aren't important.

In many circumstances one is likely to want to integrate out the second order uncertainty to produce a conventional cumulative as illustrated in

---

[5] In practice there are several important issues that can complicate such manipulations. In addition, not all curves that look like Figure 31 are equivalent. Some are "marginals" and can only be read in a vertical or a horizontal direction.

Figure 33.[6] In the case where we have constructed results such as those in Figure 31 by combining results from a large number of alternative possible physical models, this is equivalent to saying that we will assess our subjective probability of the likelihood that each alternative model describes the physical reality that we believe will actually occur, and then perform a probability weighted sum across the results.

If one is making decisions on the basis of maximized expected utility one can in general use this integrated result rather than the full set of curves without any significant change in outcome.

Problems can arise in interpreting results of the sort displayed in Figure 31 if one is using a probabilistic statement of limit, or if one is not content to make decisions on the basis of expected net utility. In this latter case the results may begin to depend significantly on the order in which the integrals are done. While the analytics can become more complex, the basic problems are not analytical in nature, they arise from the fact that because displays such as that in Figure 31 are subtle and complex, their proper interpretation can be difficult. Whether one opts for such complex displays, or chooses to simplify by integrating out some of the detail, should depend entirely on how the decision maker plans to make the necessary decision and how useful the decision maker finds the results in providing insight about the problem he or she faces.

---

[6] In performing such an integration it is important to make sure that the curves were constructed in such a way as to make the necessary integrals meaningful. See footnote 5.

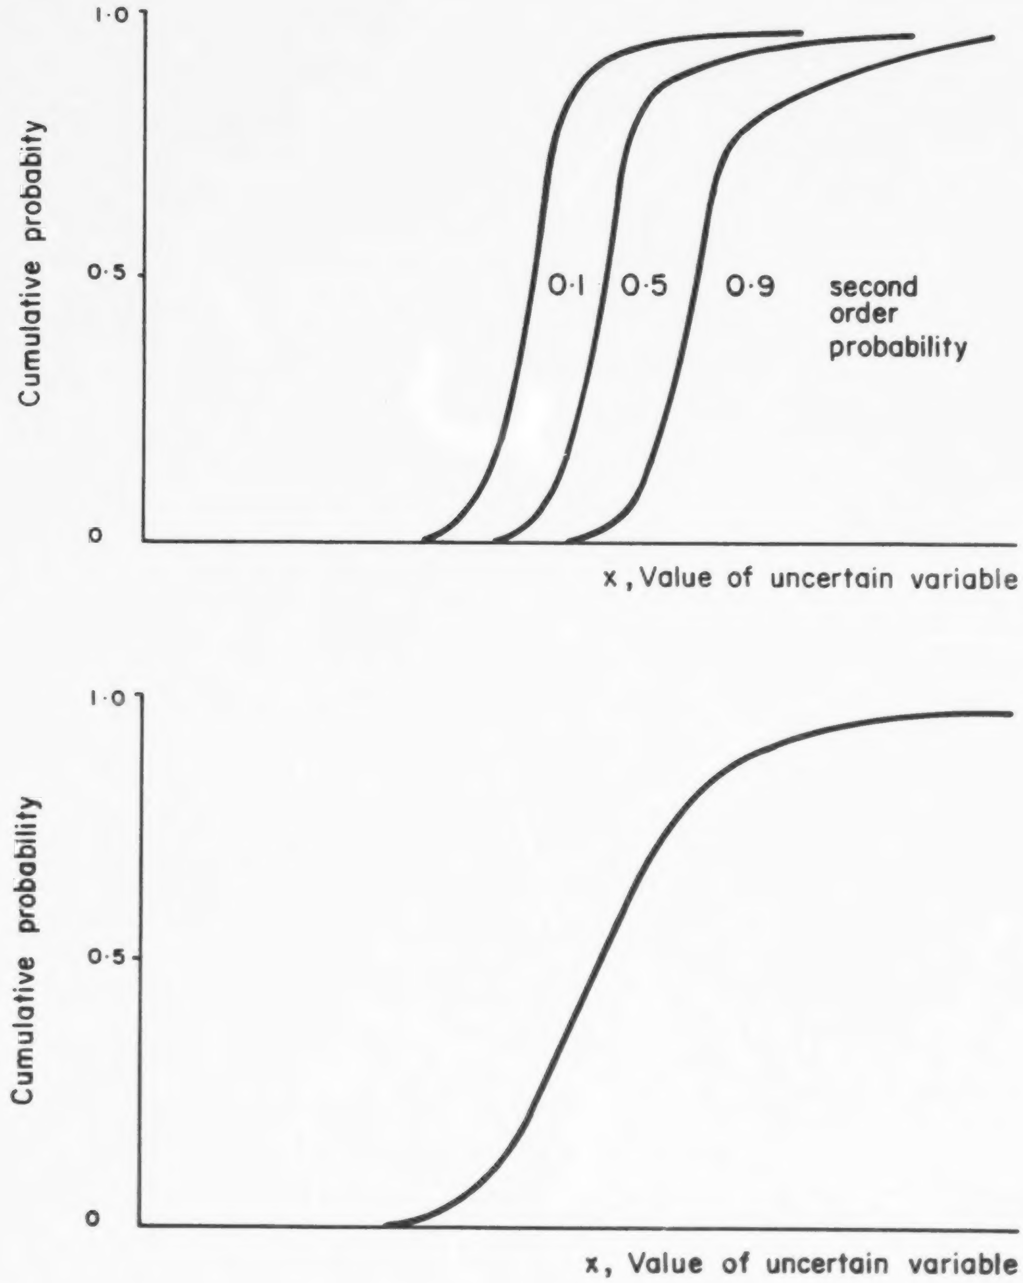

FIGURE 33. Illustration of the way in which a CDF with second order probability displayed (above) can be integrated to yield a conventional CDF (below).

Conclusion

The bottom line is quite simple. The point of doing analysis on risk management problems, such as the uranium mill tailing problem, should be to try to frame the problems we face in useful and productive ways and develop insights that will help us do a better job in risk management. The literatures of decision and policy analysis have developed a large number of analytical and display tools that can be used in this process. They are not however ends in themselves. They should be used only if they help to give insight and understanding. The choice of which to use should depend on how helpful...after suitable practice and experience... analysts and decision makers find them.

References

1. See Potential Health and Environmental Hazards of Uranium Mine Wastes, EPA 520/1-83-007, June 1983; Final Environmental Impact Statement for Standards for the Control of Byproduct Materials from Uranium Ore Processing, EPA 520/1-83-008, September 1983; and "Standards for remedial actions at inactive uranium processing sites", 40 CFR Part 192, Federal Register, 48(3) pp. 590-604.

2. M. Granger Morgan, "Some Fundamental Issues in the Development of a Coherent Risk Management Strategy for Uranium Mill Tailings", Appendix D in Approaches to Risk Assessment for Canadian Mill Tailings, DSS File 14SQ.23241-4-1695, a report from Senes Consultants Inc. to The Canadian National Uranium Mill Tailings Program, October 1985.

3. William Hogan as quoted by Max Henrion in a paper presented to the annual meeting of the Society for Risk Analysis, Knoxville, Tennessee, 1985.

4. Some of the discussion presented in this section draws upon ideas developed by the author and Max Henrion in their draft of "Chapter 11 The Graphic Communication of Uncertainty" in the book Uncertainty: a guide to the characterization and treatment of scientific and technical uncertainty in risk and other quantitative policy analysis now being written by these two in collaboration with Gregory McRae and Mitchell Small.

5. M. Granger Morgan, Samuel C. Morris, Max Henrion, Deborah A. L. Amaral, and William R. Rish, "Technical Uncertainty in Quantitative Policy Analysis: A sulfur air pollution example", Risk Analysis, 4, pp. 201-216, 1984.

6. M. Granger Morgan, Max Henrion, Samuel C. Morris, and Deborah A. L. Amaral. "Uncertainty in Risk Assessment: a case study involving sulfur air pollution", Environmental Science and Technology, 19, pp. 662-667, 1985.

7. M. Granger Morgan, "Technical Uncertainties in the Policy Making Process", approximate transcript of a luncheon address, 6th Annual Conference of the Institute of Environmental Quality, State of Illinois, Champaign, IL, 1977 October.

8. John D. Graham and James W. Vaupel, "Value of Life: what difference does it make?", Risk Analysis, 1, pp. 89-95, 1981.

9. For discussions of decision analysis and related topics see:

- Jack R. Benjamin and Allin Cornell, Probability, Statistics and Decision for Civil Engineers, McGraw-Hill, 1970.
- Samuel E. Bodily, Modern Decision Making: a guide to modeling with decision support systems, McGraw-Hill, 1984.
- Morris H. DeGroot, Optimal Statistical Decision Theory, McGraw-Hill, 1970.
- Ronald A. Howard, James E. Matheson, and Katherine L. Miller (eds.), The Principles and Applications of Decision Analysis, The Strategic Decisions Group, Menlo Park, CA, 1983.
- Ralph L. Keeney, "Decision analysis: an overview", Operations Research, 30, pp. 803-891, 1982 September/October.
- Howard Raiffa, Decision Analysis: Introductory lectures on choices under uncertainty, Addison-Wesley, 1970.
- Howard Raiffa and Robert Schlaifer, Applied Statistical Decision Theory, Harvard Business School 1961, and John W. Pratt, Howard Raiffa and Robert Schlaifer, Introduction to Statistical Decision Theory, McGraw-Hill, 1965.
- Robert L. Winkler, Introduction to Bayesian Inference and Decision, Holt, Rinehart and Winston, 1972.

10. Ralph L. Keeney and Howard Raiffa, Decisions With Multiple Objectives: preferences and value tradeoffs, John Wiley, 1976.

11. See for example C.-A.S. Stael von Holstein and James E. Matheson, A Manual for Encoding Probability Distributions, SRI International, Menlo Park California, August 1979; D. W. Boyd and S. G. Regulinski, Characterizing Uncertainty in Technology Cost and Performance, Decision Focus Inc., Los Altos, California, June 1979; Granger Morgan, Max Henrion and Sam Morris, Expert Judgments for Policy Analysis, Brookhaven National Laboratory report 51353, 1981.

12. See various discussions and references in Daniel Kahneman, Paul Slovic and Amos Tversky (eds.), Judgment Under Uncertainty: heuristics and biases, 555 pp., Cambridge University Press, 1982.